Modest set of simulations of local-type primordial non-Gaussianity in the DESI era

Boryana Hadzhiyska[1,2,*] Lehman H. Garrison,[3] Daniel J. Eisenstein,[4] and Simone Ferraro[1,2]

[1]*Physics Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*
[2]*Berkeley Center for Cosmological Physics, Department of Physics, University of California, Berkeley, California 94720, USA*
[3]*Scientific Computing Core, Flatiron Institute, 162 5th Avenue, New York, New York 10010, USA*
[4]*Center for Astrophysics, Harvard and Smithsonian, 60 Garden Street, Cambridge, Massachusetts 02138, USA*

A measurement of a primordial non-Gaussianity (PNG) signal through late- or early-Universe probes has the potential to transform our understanding of the physics of the primordial Universe. While large-scale structure observables in principle contain vital information, interpreting these measurements is challenging due to poorly understood astrophysical effects. Luckily, N-body simulations, such as the AbacusPNG set presented in this study, consisting of 9 boxes, each of size $L_{box} = 2$ Gpc/h and particle mass of $1.01 \times 10^{10} M_{\odot}/h$, provide a viable path forward. As validation, we find good agreement between the simulations and our expectations from one-loop perturbation theory (PT) and the "separate universe" method for the matter bispectrum, matter power spectrum and the halo bias parameter associated with PNG, b_{ϕ}. As a science application, we investigate the link between halo assembly bias and b_{ϕ} for halo properties known to play a vital role in accurately predicting galaxy clustering: concentration, shear (environment), and accretion rate. We find a strong response for all three parameters, suggesting that the connection between b_{ϕ} and the assembly history of halos needs to be taken into account by future PNG analyses. We further perform the first study of the b_{ϕ} parameter from fits to early DESI data of the luminous red galaxy (LRG) and quasistellar object (QSO) samples and comment on the effect on f_{NL} constraints for the allowed galaxy-halo models (note that $\sigma[f_{NL}] \propto \frac{\sigma[b_{\phi}]}{b_{\phi}}$). We find that the error on f_{NL} is 15, 8, 7 for the LRGs at $z = 0.5$ and $z = 0.8$ and QSOs at $z = 1.4$, respectively, suggesting that a thorough understanding of galaxy assembly bias is warranted so as to perform robust high-precision analysis of local-type PNG with future surveys.

I. INTRODUCTION

The investigation of the nature of primordial density fluctuations has long been a cornerstone of modern cosmology. Among the many tantalizing prospects of understanding these fluctuations, finding non-Gaussian features imprinted on them during the earliest moments after the big bang would allow us to put to the test the most widely accepted paradigms about the primordial Universe and illuminate the physics of the high-energy regime [1].

The study of primordial non-Gaussianity (PNG) encompasses a diverse array of observational avenues ranging from early- to late-Universe probes, seeking to discern subtle deviations from Gaussian statistics in the structures of our Universe. Particularly interesting for galaxy surveys is the detection and characterization of the so-called "local-type PNG," as this type of PNG links the small-scale with the large-scale galaxy density distribution. A number of analyses have tried to constrain the parameter characterizing the PNG amplitude, f_{NL}, both in the cosmic microwave background (CMB) [2] and in the large-scale structure [3].

The simplest single-field models of inflation predict vanishing f_{NL} [4–7], and hence, any detection of a nonzero signal would impact our understanding of the mechanism that generated the seeds of structure formation. The current best constraints on local PNG come from the analysis of the CMB by the *Planck* satellite, finding $f_{NL} = -0.9 \pm 5.1(1\sigma)$ [8]. The tightest near-future constraints from large-scale structure surveys are theoretically predicted to reach $\sigma(f_{NL}) \approx 1$, as the 3D distribution of galaxies from future surveys contains more Fourier modes than the 2D CMB map [9,10], contaminated by foregrounds and Silk-damping on small angular scales. While theoretically these improvements are significant, there are still difficulties associated with interpreting the observed signal. In particular, to uncover these subtle non-Gaussian imprints, we need to disentangle astrophysical and nonlinear

*boryanah@berkeley.edu

TABLE I.   AbacusPNG parameter specifications.

Name	f_{NL}	Phase	Full outputs[a]	Comment
Abacus_pngbase_c000	0	ph000	Full	Vanilla ΛCDM at pngbase mass resolution
Abacus_pngbase_c{300..303}	{30, −30, 100, −100}	ph000	Full	f_{NL} variations, otherwise c000 cosmology
Abacus_pngbase_c000	0	ph001	Partial	Second realization, fewer full particle outputs
Abacus_pngbase_c{300..303}	{30, −30, 100, −100}	ph001	Partial	Second realization, fewer full particle outputs

[a]"Full outputs" refers to the set of output epochs for which complete particle snapshots are stored (as opposed to subsamples). See Sec. II for details.

evolution effects from the features of the primordial Universe.

Summary statistics such as higher-order correlation functions, are particularly sensitive to all PNG shapes and provide a powerful venue for conducting tests on our cosmological observables. In the case of local PNG, there is a strong response even in the galaxy two-point correlation function through the scale-dependent bias. Thus, this feature, which is absent in the matter field, allows us to discern the imprints of local PNG by analizing the galaxy power spectrum on large scales. Several attempts to measure the scale-dependent bias using quasars from spectroscopic galaxy surveys have already been made, and projections for near-future experiments have forecast tighter constraints than future CMB observations. However, challenges persist in translating PNG detections into specific primordial bispectrum amplitudes, which is crucial for validating or ruling out single-field inflationary models [11,12].

In recent years, the utilization of state-of-the-art numerical simulations, both hydrodynamical and N-body, has revolutionized our ability to explore the complex interplay between primordial physics and large-scale structure e.g., [13–18]. These simulations provide an invaluable laboratory for studying the link between galaxy formation and primordial physics, which is crucial for disentangling the amplitude of PNG, parametrized by f_{NL}, from the galaxy bias response to long-wavelength fluctuations, parametrized by b_ϕ, which are otherwise degenerate in standard scale-dependent bias analyses.

Hydrodynamical simulations provide us with a full set of realistic galaxy properties, including galaxy colors, black hole mass, and stellar mass. Because they are very computationally expensive, the amount of volume that can be simulated is typically quite limited ($\lesssim 1$ Gpc^3). However, in recent years, the technique of "separate Universes" has been employed to understand the response of small-scale galaxy physics to long-wavelength fluctuations, rendering the volume limitation a lesser issue see e.g., [19]. Nonetheless, relying on a single implementation of the physical model can also be problematic and bias our understanding of the relationship between local PNG and galaxy formation. On the other hand, N-body simulations typically have a much larger volume, but since they lack galaxy physics, one needs to use heuristic methods to paint galaxy properties on them. Luckily, recently there has been significant progress in

developing high-fidelity models and applying them to large-volume N-body simulations e.g., [20–23].

In this work, we present a new set of N-body simulations run with the Abacus code [24] that incorporates local-type PNG, in conjunction with the latest galaxy-halo models applied in the analysis of current cosmological spectroscopic surveys [25] to study the connection between scale-dependent bias, b_ϕ, and the intrinsic properties of galaxies and halos. Another relevant simulation suite that features PNG is the Quijote-PNG suite, which besides local-type incorporates two other types of PNG. While its box size is a bit more modest ($L_{box} = 1$ Gpc$/h$), the suite comprises ~1000 boxes and has been used to study both the information content of the halo and the matter density field [15,16]. In Sec. III we validate the power spectrum and bispectrum against the one-loop perturbation theory (PT) prediction [26] to ensure that the theoretical prediction matches the simulations on large scales. In Sec. IV, we then study the response of b_ϕ to relevant halo properties for galaxy formation such as environment and concentration. Finally, in Sec. V, we characterize the value of b_ϕ for faithful representations of the large-scale experiment DESI and make simple Fisher forecasts for the prospects of detecting f_{NL} through scale-dependent bias of the DESI galaxy tracers. We summarize our findings in Sec. VI.

II. SIMULATIONS

A. Description

AbacusPNG consists of 10 simulations: five variations of f_{NL} with two realizations each. The key parameters of the simulations are summarized in Tables I and II. In brief, the simulations span $f_{NL} = \{-100, -30, 0, 30, 100\}$ with 4096^3 particles in boxes of $2h^{-1}$ Gpc, yielding a particle mass of $1 \times 10^{10} h^{-1} M_\odot$.

The AbacusPNG simulations are effectively an extension of the AbacusSummit simulations[1] [27]. They share the same scheme for cosmologies and phases (c000, ph000, etc) so that the two sets can be meaningfully compared. The parameters controlling the numeric accuracy of the solution are the same; likewise, the output data model is the same, including cleaned CompaSO halo catalogs [28,29]),

---

[1]abacussummit.readthedocs.io/.

TABLE II. Common simulation parameters for the AbacusPNG set of simulations.

Parameter	Value	Comment
N	4096^3	Number of particles
L	$2000h^{-1}$ Mpc	Box size
z_{init}	99	Initial redshift
z_{final}	0.3	Final redshift
ϵ	$12.2h^{-1}$ kpc	Proper softening length, Plummer-equivalent
η_{acc}	0.25	Time step parameter, acceleration-based
M_{p}	$1 \times 10^{10}h^{-1}M_{\odot}$	Particle mass

light cones, power spectra, merger trees, and particle subsamples. This shared data model means that the ABACUSUTILS interface[2] can be used to interact with both sets of simulations.

Despite their similarities, the base mass resolution of AbacusPNG is 4.8× coarser than that of AbacusSummit, which is $2.1 \times 10^9 \, h^{-1}M_{\odot}$. And AbacusSummit, while containing several variations in particle mass, does not include the particular AbacusPNG resolution. Therefore, AbacusPNG includes two vanilla (non-PNG) ΛCDM simulations, `Abacus_pngbase_c000_ph{000,001}`, that can be compared to AbacusSummit to isolate the effects of mass resolution from the effects of PNG.

The low mass resolution (compared to AbacusSummit) is by design, as we expect most information on PNG in DESI to come from LRGs and QSOs [see Section 2.5.1 in [30]], which have higher host halo mass than ELGs, for which AbacusSummit's resolution was optimized. Therefore, at fixed computational budget, we prioritized more realizations and greater volume rather than greater resolution. This is also why the final redshift is 0.3, as the great majority of DESI-targeted LRGs and QSOs are at higher redshifts.

A very similar version of Abacus was used to run AbacusPNG as AbacusSummit, except for minor improvements to accuracy of the near-field force (better numerical stability in the accumulation of partial forces); improvements to the numerical stability of the on-the-fly 2LPT scheme (more accurate representation of displacements); and various performance optimizations for Perlmutter. Abacus was run with the same parameters controlling the numerical accuracy (in particular, multipole order and time step parameter) as AbacusSummit, and therefore the limits on accuracy and convergence placed by [24,27] can be conservatively applied to AbacusPNG, too.

The output redshifts of various data products are as follows:

(i) CompaSO halo catalogs (29 redshifts): 8.0, 5.0, 3.0, 2.75, 2.5, 2.25, 2.0, 1.85, 1.7, 1.625, 1.55, 1.475, 1.4, 1.325, 1.25, 1.175, 1.1, 1.025, 0.95, 0.875, 0.8, 0.725, 0.65, 0.575, 0.5, 0.45, 0.4, 0.35, 0.3.

(ii) Particle subsamples, 3% and 7% sets (10 redshifts): 3.0, 2.5, 2.0, 1.7, 1.4, 1.1, 0.8, 0.5, 0.4, 0.3.

(iii) Full outputs, "Partial" list (7 redshifts): 3.0, 2.5, 2.0, 1.4, 0.8, 0.5, 0.3.

(iv) Full outputs, "Full" list (10 redshifts): 3.0, 2.5, 2.0, 1.7, 1.4, 1.1, 0.8, 0.5, 0.4, 0.3.

(v) Light cones: shells are output at every time step, approximately 800 epochs with 3 "observers" (same configuration as AbacusSummit).

The `ph000` simulations use the "Full" list of full outputs, and `ph001` uses the "Partial" list.

B. PNG implementation

Local-type PNG is parameterized in terms of the primordial gravitational potential during matter domination, $\phi(\mathbf{x})$, and the parameter f_{NL}, which quantifies the amount of non-Gaussianity, via [31]:

$$\phi(\mathbf{x}) = \phi_{\mathrm{G}}(\mathbf{x}) + f_{\mathrm{NL}}[\phi_{\mathrm{G}}(\mathbf{x})^2 - \langle\phi_{\mathrm{G}}(\mathbf{x})^2\rangle], \qquad (1)$$

where ϕ_{G} is a Gaussian random field.

To generate initial conditions with local-type PNG, we follow the steps outlined below:

(i) From the input power spectrum, we obtain a realization of the Gaussian primordial potential field, $\phi_{\mathrm{G}}(\mathbf{k})$.

(ii) We execute an inverse fast Fourier transform (FFT) to convert the field to real space, then square it and subtract the mean: $\phi_{\mathrm{G}}(\mathbf{x})^2 - \langle\phi_{\mathrm{G}}(\mathbf{x})^2\rangle$.

(iii) Finally, we normalize it by the desired amplitude of local-type PNG, f_{NL}, add it back to $\phi(\mathbf{x})$, and execute a forward FFT to generate $\phi(\mathbf{k})$.

The resulting field becomes the potential source term used when running the initial conditions generator. This procedure is implemented in the zeldovich-PLT code[3] [32].

In detail, to obtain the primordial potential from the input power spectrum in the first step, the code accepts the spectral index of the primordial power spectrum. The transfer function is inferred by assuming $T(k) = 1$ on large scales. This has the benefit of not requiring a separate input file, and suffices for our setup where this assumption holds true.

These initial conditions are generated at $z_{\mathrm{init}} = 99$ using an identical procedure to AbacusSummit. The initial power spectrum is that of CDM and baryons at $z = 1$, backscaled to z_{init} using the linear growth factor including a nonclustering neutrino approximation. We account for the deleterious effects of particle discreteness on the small-scale growth by applying the Particle Linear Theory rescaling of [32], with a target redshift of $z = 12$. The second-order Lagrangian perturbation theory (2LPT)

---

method from that same work is applied here, via direct evaluation of the forces in the first two Abacus time steps.

C. Deployment of Abacus on Perlmutter

The following subsection is a brief report on some technical challenges and successes in running AbacusPNG on Perlmutter, including discussion of performance. It can be safely skipped by readers only interested in the scientific aspects of this work.

AbacusPNG is the first set of Abacus simulations run on NERSC's Perlmutter,[4] an HPE Cray Shasta system. The simulations were run on Perlmutter's GPU nodes using DESI's compute allocation. Each simulation used 32 GPU nodes, each with $4 \times$ NVIDIA A100 (PCIe 4.0) GPU; $1 \times$ AMD EPYC 7763 (1 socket, 64 cores); 256 GB DDR4 RAM; and $4 \times$ HPE Slingshot 11 NICs. We used the Cray GNU programming environment with Cray FFTW and Cray MPICH (not using any GPU-aware features or GPU-GPU communication), and CUDA 11. We will highlight some technical and performance aspects of AbacusPNG here but refer the reader to our more detailed reports in [27,33] for context on Abacus internals.

As in AbacusSummit, we used a 1D toroidal "slab" parallelization scheme to distribute particles across nodes. While Abacus is now capable of 2D domain decompositions, the main benefit of 2D for Abacus is to enable larger simulations where domain width per node would otherwise become too narrow. This was not required for these simulations, so we employed the 1D strategy, which has lower communication and synchronization overheads.

Abacus performed well on Perlmutter, despite far less tuning effort being devoted to it than for AbacusSummit. Each simulation took about 13 hours of wall clock time—about 425 node-hours. Each simulation's compute rate began around 50 M particle / sec per node at z_{init} (`singlestep` and `convolution` combined) and actually maintained that rate until the final redshift of 0.3 (about 800 time steps). This rate can be compared with 70 M particle / sec per node on Summit at the initial time, which decreased to 45 M at late times. Fairly substantial Perlmutter overheads in launching each time step as a separate executable invocation degraded this performance to a mean of 36 M particle / sec. The overheads arose from variable performance of the NERSC file systems and Slurm job scheduler; the slowdowns were bursty and time-correlated, probably due to load on the system from unrelated jobs. Future upgrades to Abacus will avoid these overheads by running multiple time steps within a single executable invocation.

The fact that the simulations did not slow down as the particles evolved to a clustered state means that the time in the near-field computation did not exceed that of the far-field even in the simulations' most clustered state. Indeed, at the initial time the GPU overlapped 50% of the CPU

work, while at the final time it only overlapped 75%. The near-field work increased by $1.9\times$, but the GPU performance also increased by about $1.3\times$ (such an efficiency increase, albeit larger, was seen in AbacusSummit as well).

Having anticipated that, compared to AbacusSummit, AbacusPNG's lower mass resolution, higher terminal redshift, and Perlmutter's large GPU-to-CPU compute ratio would result in less near-field work, we opted for a low cells-per-dimension (CPD) value of 875, shifting work from the far-field (CPU) to the near-field (GPU). At 102 particles-per-cell, this is about 50% more particles per cell and $2.3\times$ more near-field work in the initial, unclustered state than AbacusSummit. Yet the $1.9\times$ growth in the near-field work was so mild that we probably could have achieved a faster time-to-solution with an even lower CPD. This mildness is attributable to the relatively large near-field/far-field transition radius in this configuration—about $5h^{-1}$ Mpc. Above the halo scale, the growth in the integrated 3D correlation function (that is, the growth in total number of pairs) falls off quite rapidly.

Finally, one technical hurdle to running on Perlmutter was an issue in restoring checkpoints to nodes using the shared memory scheme of [34]. The restore process simply consisted of copying files from Perlmutter Scratch to local storage using Python's SHUTIL package, but such copies would hang with high frequency. The root cause was found to be an issue with the implementation of the `sendfile` Linux syscall in the Lustre filesystem used by Perlmutter Scratch. `sendfile` is designed for high-performance, in-kernel copies, and issues with it on network filesystems are not uncommon. We identified the issue during the Perlmutter acceptance testing period and reported it to NERSC, along with the workaround for SHUTIL, which is to set `shutil._USE_CP_SENDFILE = False`.

III. VALIDATION

In this section, we test the output of the AbacusPNG simulations against expectations of the matter and halo field distributions, using theoretical predictions from PT and heuristic methods such as the "separate universe" approach.

A. Matter field

To validate the simulated matter field, we examine two summary statistics: the matter power spectrum and the matter bispectrum.

1. Bispectrum

We start our investigation with the three-point statistic, the matter bispectrum, which we compare against the tree-level prediction at the initial redshift of the simulation ($z_{\text{IC}} = 99$), when the Universe was matter-dominated. We note that at such high redshifts the tree-level prediction should perform well down to small scales. As the Universe

becomes more nonlinear, the tree-level approximation starts to break down, and one needs to employ a higher-order perturbative expansion [see e.g., [26]].

The matter bispectrum is defined as

$$\langle \delta_m(\mathbf{k}_1)\delta_m(\mathbf{k}_2)\delta_m(\mathbf{k}_3)\rangle$$
$$= (2\pi)^3\delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3)B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3). \quad (2)$$

We estimate it using the PYLIANS3 package, which is based on the implementation of Foreman *et al.* [35], and use linear bins between 0 and $k_{\max} = 0.4\ h/\text{Mpc}$, of width $2k_F$, where k_F is the fundamental mode, $k_F \equiv 2\pi/L_{\text{box}}$, and L_{box} is the box size of the simulation.

While there are other theoretically motivated shapes that are interesting to consider, here we restrict our analysis only to the local-type PNG [see e.g. [36,37], for reviews of PNG]. In that case, we can express the primordial bispectrum, $B_\phi(k_1, k_2, k_3)$, as

$$B_\phi(k_1, k_2, k_3) = 2f_{\text{NL}}P_\phi(k_1)P_\phi(k_2) + 2\,\text{perm.} \quad (3)$$

where $P_\phi(k)$ is the primordial power spectrum and f_{NL} is the amplitude of non-Gaussianity, associated with this shape. Taking into account the relationship between matter density and the primordial (Bardeen) potential:

$$\delta_m(\mathbf{k}, z) = \mathcal{M}(k, z)\phi(\mathbf{k}), \quad (4)$$

where

$$\mathcal{M}(k, z) = \frac{2}{3}\frac{k^2 T_m(k)}{\Omega_m H_0^2}D(z), \quad (5)$$

T_m is the matter transfer function, D is the growth rate normalized to be equal to the scale factor, a, during matter domination (i.e., CMB convention), Ω_m is the fractional matter density parameter today, and H_0 is the Hubble expansion rate today, one can convert the primordial bispectrum into the matter bispectrum via

$$B(k_1, k_2, k_3) = B_\phi(k_1, k_2, k_3)\mathcal{M}(k_1, z)\mathcal{M}(k_2, z)\mathcal{M}(k_3, z). \quad (6)$$

Local-type PNG is a powerful probe of inflation. In single-field slow-roll inflationary models, the amplitude of the bispectrum is $\ll \mathcal{O}(\eta, \epsilon)$, where ϵ and η are the slow-roll parameters Maldacena [38], Creminelli and Zaldarriaga [39]. Measuring a large signature of local-type PNG would thus rule out slow-roll, single-field inflation. Observable levels of local-type PNG are predicted by a number of multifield inflationary models, such as the curvaton and modulated reheating models [40,41], as well as by some noninflationary models [42].

In comparing simulations with theory, we confine our study to the bispectrum derivative rather than the full bispectrum, so as to isolate only the non-Gaussian contribution to the full bispectrum. In particular, we compute

$$\frac{\partial B(k_1, k_2, k_3)}{\partial f_{\text{NL}}} = 2(P_\phi(k_1)P_\phi(k_2) + 2\ \text{perm.})$$
$$\mathcal{M}(k_1, z)\mathcal{M}(k_2, z)\mathcal{M}(k_3, z). \quad (7)$$

We note that there are several subtleties in this calculation. To reduce the noise associated with the conversion from B to B_ϕ, we obtain the Bardeen potential by dividing the matter density field $\delta_m(\mathbf{k})$ by $\mathcal{M}(\mathbf{k}, z)$ evaluated on the 3D Fourier grid. In addition, when calculating the theoretical prediction for the bispectrum derivative, we first evaluate the power spectrum $P_\phi(\mathbf{k})$ on the 3D grid, then take the product of the power spectra, and finally average over the triangle configurations of interest that satisfy the Dirac delta condition, $\delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3)$. This is necessary, as the theoretical measurements calls for evaluating the average of the product rather than the product of two average quantities, i.e., $\langle ab\rangle \neq \langle a\rangle\langle b\rangle$. We find that not accounting for this effect leads to 10% differences between theory and simulations when predicting small-angle triangle shapes such as the squeezed limit bispectrum.

In Fig. 1, we show three different shapes for the local-type PNG matter bispectrum derivative at $z_{\text{IC}} = 99$, calculated by finite-differencing the AbacusPNG simulation AbacusPNG_c302_ph000 ($f_{\text{NL}} = 100$). We find very agreement between theory using the tree-level approximation [see Eq. (6)] and simulations within 5% for all three triangle configurations considered in this study: equilateral ($k_1 = k_2 = k_3$), squeezed [$k_1 = k_2 = k$, $k_3 = 3k_F$, folded ($k_1 = k_2 = k$, $k_3 = 2k$)]. The agreement is best for the equilateral case for which grid effects play the smallest role. In the case of local-type PNG, this is also the shape that contributes the least to the total information. On the other hand, we found that both the squeezed and folded shapes are affected by grid effects (see previous paragraph). Nonetheless, we find satisfactory agreement for both shapes. In fact, we have separately conducted a test against 2LPTPNG's implementation of local-type PNG and have found better agreement between the two codes than with the tree-level expression. As expected, the squeezed limit yields the strongest response to local-type PNG (across all k-modes). For $k \gtrsim 0.15$, we see a deviation from theory, which we attribute to mild nonlinearities we find in the simulation power spectrum on these scales, compared with the linear theory output.

2. Power spectrum

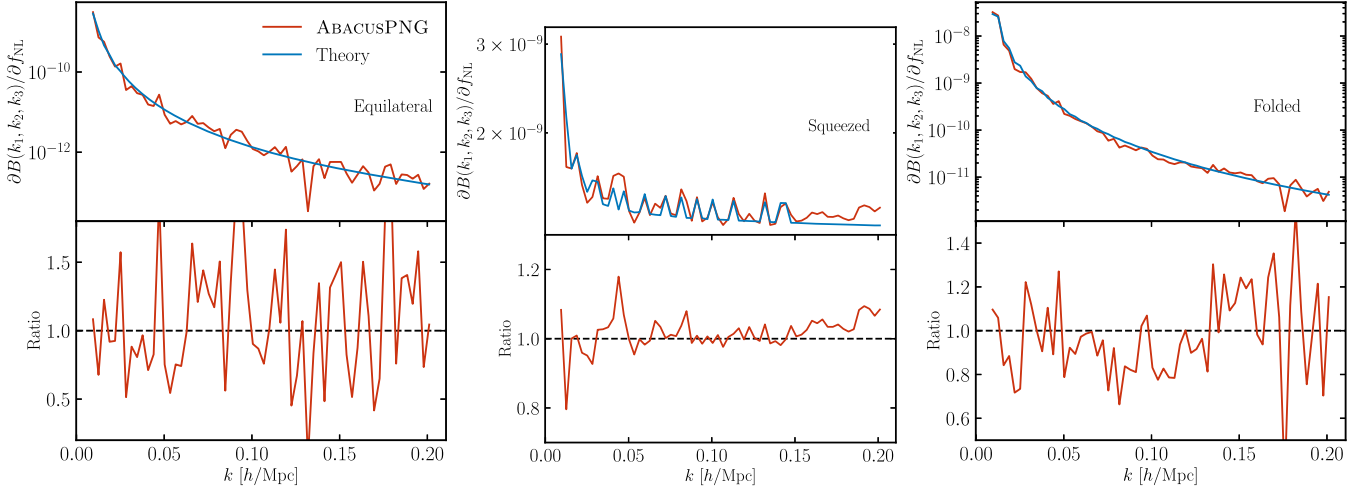Next, we consider the matter power spectrum, which is defined as

FIG. 1.   Validation of the effect of local-type PNG on the matter bispectrum at the initial conditions, $z_{IC} = 99$. We show the derivative of the power spectrum with respect to local-type f_{NL} as computed from the `AbacusPNG_c302_ph000` simulation as well as from theory using the tree-level approximation [see Eq. (6)]. We see that they are in very good agreement with each other for all three triangle configurations considered in this study: equilateral ($k_1 = k_2 = k_3$), squeezed ($k_1 = k_2 = k$, $k_3 = 3k_F$), folded ($k_1 = k_2 = k$, $k_3 = 2k$). The squeezed limit yields the strongest response to local-type PNG (across all k-modes) and thus has the smallest error bars. For $k \gtrsim 0.15$, we see a deviation from theory, which we attribute to mild nonlinearities we find in the simulation power spectrum on these scales, compared with linear theory. The noise in the theory curve is due to the fact it is computed on a grid so as to match the noise of the measurement.

$$\langle \delta_m(\mathbf{k})\delta_m^*(\mathbf{k'})\rangle = (2\pi)^3\delta^{(3)}(\mathbf{k} - \mathbf{k'})P_{mm}(k). \quad (8)$$

To compute the power spectrum in the AbacusPNG simulations for the LRGs (QSOs), we use the ABACUSUTILS code with linear bins of width $0.05\ h/\text{Mpc}$, from 0 to $k_{max} = 0.4(0.32)h/\text{Mpc}$.

Since at the tree-level, the PNG contribution to the matter power vanishes (we will later see that in the case of galaxy tracers, that is not the case), we need to go to higher order in order to obtain a theoretical prediction for the matter power spectrum. In particular, we adopt the 1-loop PT expression from Cabass *et al.* [26], as follows:

$$P(\mathbf{k}) = P_G(\mathbf{k}) + f_{NL}P_{12}(\mathbf{k}), \quad (9)$$

where P_G is the standard Gaussian contribution to the power spectrum, and the PNG contribution can be calculated by solving the integral:

$$P_{12}(\mathbf{k}) = 2\int \frac{\mathrm{d}^3q}{(2\pi)^3} F_2(\mathbf{q}, \mathbf{k} - \mathbf{q})B(k, q, |\mathbf{k} - \mathbf{q}|) \quad (10)$$

where B is the matter power spectrum [see Eq. (6)] and F_2 is the standard perturbation theory (SPT) kernel defined as:

$$F_2(\mathbf{k}_1, \mathbf{k}_2) = \frac{5}{7} + \frac{2}{7}\frac{(\mathbf{k}_1 \cdot \mathbf{k}_2)^2}{k_1^2 k_2^2} + \frac{1}{2}\frac{\mathbf{k}_1 \cdot \mathbf{k}_2}{k_1 k_2}\left(\frac{k_1}{k_2} + \frac{k_2}{k_1}\right). \quad (11)$$

We are only interested in the derivative with respect to f_{NL}, which becomes:

$$\frac{\partial P(\mathbf{k})}{\partial f_{NL}} = 2\int \frac{\mathrm{d}^3q}{(2\pi)^3} F_2(\mathbf{q}, \mathbf{k} - \mathbf{q})B(k, q, |\mathbf{k} - \mathbf{q}|) \quad (12)$$

To obtain the derivative, we perform Vegas Monte Carlo integration using PyCUBA from the PyMultiNest package.[5]

Our results at $z = 0.5$ are shown in Fig. 2. We see that the theory prediction of the derivative matches the simulation result very well until $k \approx 0.1\ h/\text{Mpc}$, after which the 1-loop approximation starts to break down. As can be seen from the figure, the signature on the matter power spectrum is extremely small ($\ll 0.1\%$) for all relevant scales ($k \lesssim 0.5\ h/\text{Mpc}$). This can be explained by the fact that unlike the bispectrum, f_{NL} does not contribute at the tree-level, so one has to go to 1-loop to see an effect. On the scales where theory and simulations approximately agree, $k \approx 0.2\ h/\text{Mpc}$, the fractional change due to PNG, $\partial \ln P_{mm}/\partial f_{NL}$ becomes 0.0001. We attribute the remaining difference between the two to higher-order contributions. From the figure, we also see that the effect on large scales is negligible, whereas on small scales it grows exponentially. As expected, the agreement between theory and simulation is very good on large scales and starts to break down on nonlinear scales $k \gtrsim 0.15\ h/\text{Mpc}$.

B. Halo field

The distribution of halos and galaxies is a potent probe of local-type PNG via the scale-dependent bias feature in the power spectrum [e.g., [3,43–45]]. In particular,

---

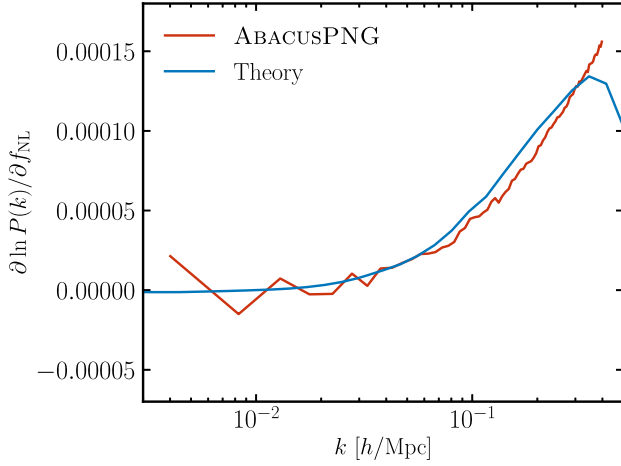[5]https://github.com/JohannesBuchner/PyMultiNest.

FIG. 2.   Validation of the effect of local-type PNG on the matter power spectrum at $z = 0.5$. We show the derivative of the power spectrum with respect to f_{NL} as computed from the `AbacusPNG_c302_ph000` simulation as well as from theory, using 1-loop EFT [see Eq. (12)]. We notice that local-type PNG has a very weak effect on the matter power spectrum – $\sim 0.01\%$ at $k \approx 0.2$ h/Mpc The effect on large scales is negligible, whereas on small scales it grows exponentially. As expected, the agreement between theory and simulation is very good on large scales and starts to break down on nonlinear scales $k \gtrsim 0.15$ h/Mpc, where nonlinearities start to become relevant.

in the presence of PNG, the galaxy/halo bias expansion becomes

$$\delta_h(\mathbf{x}, z) = b_1(z)\delta_m(\mathbf{x}, z) + b_\phi(z)f_{\mathrm{NL}}\phi(\mathbf{x}) + \epsilon(\mathbf{x}), \quad (13)$$

where $\delta_h(\mathbf{x}, z) = n_h(\mathbf{x}, z)/\bar{n}_h(z) - 1$ with n_h as the local number density of galaxies/halos. Similarly to the matter power spectrum, which we defined in Sec. III A, we can define the galaxy/halo power spectrum as:

$$\langle \delta_h(\mathbf{k})\delta_h(\mathbf{k}')\rangle = (2\pi)^3\delta^{(3)}(\mathbf{k} - \mathbf{k}')P_{hh}(k) \quad (14)$$

To lowest order, we can express it as:

$$\begin{aligned}
P_{hh}(k, z) &= b_1^2 P_{mm}(k, z) + 2b_1 b_\phi f_{\mathrm{NL}} P_{m\phi}(k, z) \\
&\quad + b_\phi^2 f_{\mathrm{NL}}^2 P_{\phi\phi}(k) + P_{\epsilon\epsilon}(k) \\
&= \left[b_1 + \frac{b_\phi f_{\mathrm{NL}}}{\mathcal{M}(k, z)}\right]^2 P_{mm}(k, z) + P_{\epsilon\epsilon}, \quad (15)
\end{aligned}$$

where $\mathcal{M}$ is defined in Eq. (5), $P_{m\phi}$ denotes the cross-power spectrum between the primordial potential and the matter field, and $P_{\epsilon\epsilon}$ is the noise power spectrum, which we can assume to be scale-independent. Note that the cross-power term is smaller than the autoterm for $k \to 0$. We can see this by noticing that on large scales ($k \lesssim 0.01$ h/Mpc), the presence of PNG induces scale-dependent corrections $\propto b_1 b_\phi f_{\mathrm{NL}}/k^2$ and $\propto b_\phi^2 f_{\mathrm{NL}}^2/k^4$ relative to P_{mm}. As we

push down to larger and larger scales, the k^{-4} dominates the scale-dependent signal (the transfer function is approximately equal to 1 on these scales). However, the amplitude of this effect is determined by the product $b_\phi f_{\mathrm{NL}}$, rendering searches for PNG using the scale-dependent bias critically dependent on our ability to determine the bias parameter, b_ϕ.

The bias parameter, b_ϕ, can be estimated assuming universality of the halo mass function as [46–53]:

$$b_\phi(z) = 2\delta_c(b_1(z) - 1), \quad (16)$$

where $\delta_c = 1.686$ is the threshold overdensity for spherical collapse. Note that this is the value one obtains from linearly extrapolating to $z = 0$. When comparing this relation to N-body simulations, recent works have found that the universality relation provides a good approximation for $b_1 \lesssim 1.5$, but overpredicts the numerical results for larger biases [13,52–59]. We note that this relation also depends on the definition of halo mass definition. A commonly adopted modification takes the following form:

$$b_\phi(z) = 2c\delta_c(b_1(z) - p), \quad (17)$$

with $c \approx 0.9$ and $p = 1$, 1.6, where the latter value is adopted for quasars, as they are believed to be "recently accreted" [46], whereas the former value is typically adopted for magnitude-limited and red galaxy samples.

This $b_\phi(b_1)$ relation is employed when placing constraints and making forecasts on local-type f_{NL} from the galaxy power spectrum [46,60–75]. Since the two quantities are degenerate in the galaxy power spectrum, it is evident that uncertainties on our theoretical modeling of b_ϕ translate directly into uncertainties on our f_{NL} measurements. Later in this paper, we explore how the $b_\phi(b_1)$ relation changes for halos selected by different intrinsic property and also for galaxies targeted by the ongoing cosmological survey DESI.

C. Separate universe approach

In this section, we describe the separate universe technique, which allows us to predict the galaxy bias, b_1, and the local-type PNG-induced bias, b_ϕ, by invoking the equivalence between the response of galaxy formation to long-wavelength perturbations and the response of galaxy formation to changes in the background cosmology. Assuming that the physics of galaxy formation acts on much smaller scales relative to the size of the long-wavelength perturbations, these long-wavelength perturbations act as a modified background to the process of galaxy formation on small scales [this is known as the "peak-background split" argument; see e.g., [76,77]]. In other words, the formation of halos and galaxies at fixed cosmology in some region of space embedded in a long-wavelength fluctuation, is equivalent to the formation of

galaxies and halos in a modified cosmology at cosmic mean. This constitutes the separate universe argument.

For measuring the linear bias b_1, the modified cosmology needs to have a different background matter density, whereas for measuring b_ϕ, the modified cosmology needs to have a different amplitude of the primordial scalar power spectrum, A_s, or equivalently, the amplitude of the linear power spectrum on the scale of 8 Mpc/h. In this study, we obtain the linear bias, b_1, without the use of the separate universe technique, by fitting b_1 and A in the power spectrum ratio:

$$\frac{P_{gm}(k)}{P_{mm}(k)} = b_1 + Ak^2, \qquad (18)$$

as we include the lowest-order nonlinear bias, which is proportional to k^2. We can similarly obtain the bias from the ratio of the autopower spectrum and find that the two methods are generally in agreement within 3%. Throughout the paper, we opt to quote the value of b_1 coming from the cross-power spectrum, as it does not contain a shot noise contribution, which if disregarded could bias our estimate of the bias.

In the presence of local-type PNG, the amplitude of the small-scale primordial power spectrum, $P_{\phi\phi}$, gets modulated by a long-wavelength perturbation of the primordial gravitational potential, $\phi(\mathbf{x})$, which impacts the formation of structure at the scale of the small-scale perturbations. In other words, local-type PNG induces a nonvanishing bispectrum in $\phi(\mathbf{x})$, which peaks in the so-called "squeezed limit." In the squeezed limit, two of the three legs have large values in Fourier space (i.e., short-scale), whereas the third leg has a small value of k (i.e., large-scale). Physically, this means that there is a large coupling between the long-wavelength perturbations of the primordial potential, $\phi(\mathbf{x})$, with the power spectrum of two short-scale modes, k_{short}. In other words, the primordial power spectrum at some point in space, $\mathbf{x}$, can be written as:

$$P_{\phi\phi}(k_{\text{short}}, z|\mathbf{x}) = P_{\phi\phi}(k_{\text{short}}, z)[1 + 4f_{\text{NL}}\phi(\mathbf{x})]. \qquad (19)$$

Thus, galaxies embedded in a long-wavelength perturbation "see" locally only a spatially uniform change to the variance of the fluctuations. They form as though in a separate universe with a modified amplitude of the primordial fluctuations [3,46]

$$\tilde{A}_s = A_s[1 + \delta A_s], \quad \text{with} \quad \delta A_s = 4f_{\text{NL}}\phi_L, \qquad (20)$$

where ϕ_L is to be treated as a constant locally and denotes the amplitude of the long-wavelength potential perturbation.

Mathematically, the PNG-induced bias, b_ϕ, is defined as

$$b_\phi \equiv \frac{\mathrm{d}\ln n_h(z)}{\mathrm{d}f_{\text{NL}}\phi}, \qquad (21)$$

which using Eq. (20), can be expressed as:

$$b_\phi = 4\frac{\mathrm{d}\ln n_h(z)}{\mathrm{d}\delta A_s} = 2\frac{\mathrm{d}\ln n_h(z)}{\mathrm{d}\delta\sigma_8} \qquad (22)$$

Thus, we can evaluate b_ϕ in a separate universe with a different value of A_s (σ_8). In practice, when working with cosmological simulations, this amounts to generating simulations with the same initial seed as the fiducial box, but with an input power spectrum file multiplied by $[1 + \delta\sigma_8^2]$. In observations, making this measurement directly is extremely challenging, but recent works have suggested a possible path forward that directly extracts the PNG signal from the galaxy density field [78].

In this work, we employ the "linear derivative" AbacusSummit boxes, `base_c112_ph000` and `base_c113_ph000`, which have the same initial seed as the fiducial simulation `base_c000_ph000`, but a different value of σ_8: namely, $\sigma_8^{\text{high}} = 1.02 \times \sigma_8^{\text{fid}}$ and $\sigma_8^{\text{low}} = 0.98 \times \sigma_8^{\text{fid}}$, respectively.

We estimate b_ϕ in the separate universe approach as

$$b_\phi(z) = \frac{b_\phi^{\text{high}}(z) + b_\phi^{\text{low}}(z)}{2}, \qquad (23)$$

where

$$b_\phi^{\text{high}}(z) = \frac{2}{\delta\sigma_8^{\text{high}}}\left[\frac{n_h^{\text{high}}(z)}{n_h^{\text{fid}}(z)} - 1\right], \qquad (24)$$

$$b_\phi^{\text{low}}(z) = \frac{2}{\delta\sigma_8^{\text{low}}}\left[\frac{n_h^{\text{low}}(z)}{n_h^{\text{fid}}(z)} - 1\right], \qquad (25)$$

where $n_h^{\text{fid}}(z)$, $n_h^{\text{high}}(z)$ and $n_h^{\text{low}}(z)$ is the number density of halos (galaxies) in the fiducial, high- and low-σ_8 simulations at some redshift z. Note that in this work, we are interested in studying the bias b_ϕ in different bins of halo properties (mass, concentration). Thus, as long as we hold fixed the mass/concentration/etc. cuts fixed for all three simulations, we can estimate b_ϕ in each bin of some (or more than one) halo property.

D. Comparing AbacusPNG to separate universe

As a validation of the halo catalogs of the AbacusPNG simulation set, we compare the inferred b_ϕ using the separate universe approach to b_ϕ obtained by fitting the scale-dependent bias in the halo power spectrum. This is complementary to validating the matter field through the bispectrum and power spectrum, as it allows us to test whether the local-type PNG implemented into the

simulation correctly couples the long wavelength mode with the short scales on which halo and galaxy formation takes place, resulting in a scale-dependent bias.

In particular, to obtain b_ϕ in the AbacusPNG simulations, we perform fits to the ratio between the halo power spectrum with $f_{\rm NL} \neq 0$ and the halo power spectrum with $f_{\rm NL} = 0$ for pairs of simulations with the same initial seed. That way, the cosmic variance and noise associated with the halo field cancel, and the resulting ratio yields an accurate estimate of b_ϕ. The ratio used in the fitting of b_ϕ is the following:

$$\frac{P_{hh}^{f_{\rm NL}=100}(k)}{P_{hh}^{f_{\rm NL}=0}(k)} = \left[\frac{b_1 + \frac{b_\phi f_{\rm NL}}{\mathcal{M}(k,z)}}{b_1}\right]^2, \qquad (26)$$

We note that although this is not a feasible way of measuring b_ϕ that could be adopted in observations, it is useful nonetheless for the purpose of testing our simulation outputs. Here, we employ the `AbacusPNG_c302_ph000` ($f_{\rm NL} = 100$) and the `AbacusPNG_c000_ph000` ($f_{\rm NL} = 0$) boxes. We split the halos into 12 logarithmic mass bins ranging between 10^{12} and $10^{14.4} M_\odot/h$ with the last bin encompassing all halos above $10^{14.4} M_\odot/h$. We simultaneously fit b_1 and b_ϕ from this ratio. In contrast, for the separate universe method, we obtain b_1 from Eq. (18) and b_ϕ from Eq. (23).

We show this comparison in Fig. 3 by fitting the power spectrum ratio [see Eq. (26)] with AbacusPNG and b_ϕ obtained from the separate universe approach by using a pair of the AbacusSummit suite of simulations [see Eq. (23)]. We note that the PNG result is obtained by fitting the ratio separately for `AbacusPNG_c302_ph000` and `AbacusPNG_c302_ph001` (with their respective $f_{\rm NL} = 0$ simulations) and then averaging the two values of b_ϕ. Since we only have two realizations, we would not benefit from including the $f_{\rm NL} = 30$ case. The halos are split into 12 logarithmic mass bins ranging from 10^{12} to $10^{14.4} M_\odot/h$, and the biases are computed for each bin by adopting the empirical relation between halo mass and linear bias from Tinker *et al.* [79]. We also show in blue the curve coming from the modified universality relation [see Eq. (17)]. We see that the halos in the separate universe approach are in perfect agreement with the modified universality relation. On the other hand, the AbacusPNG curve is substantially more noisy, as despite the fact that we have canceled most of the cosmic variance, the power spectrum ratio retains some intrinsic noise. Some of that noise is due to higher-order contributions, which we have ignored in the ratio [such as the coupling between δ_h and ϕ via the b_ϕ, δ parameter, see [80], for details], and some of it is due to poor statistics, which affects the high mass end, where we need to compute the power spectrum for $\sim$100 s of halos.
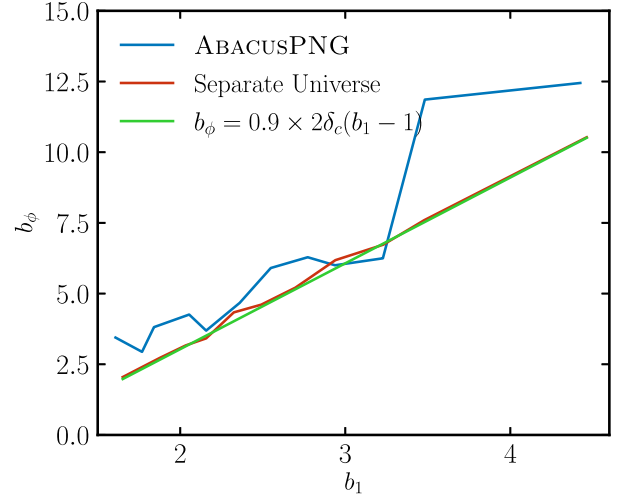

FIG. 3. Comparison between b_ϕ obtained from the AbacusPNG set by fitting the power spectrum ratio [see Eq. (26)] and b_ϕ obtained from the separate universe approach by using a pair of the AbacusSummit suite of simulations [see Eq. (23)]. We note that the PNG result is obtained by fitting the ratio separately for `AbacusPNG_c302_ph000` and `AbacusPNG_c302_ph001` (with their respective $f_{\rm NL} = 0$ simulations) and then averaging the two values of b_ϕ. We also show in blue the curve coming from the modified universality relation [see Eq. (17)]. We see that the halos in the separate universe approach are in perfect agreement with the modified universality relation. On the other hand, the AbacusPNG curve is substantially more noisy, as despite the fact that we have canceled most of the cosmic variance, the power spectrum ratio retains some intrinsic noise.

IV. HALOS

In this section, we study the relationship between the PNG-induced bias, b_ϕ, and halo assembly bias properties beyond its mass (linear bias). Halo assembly bias is defined as the response of the two-point clustering of halos to the values of intrinsic halo properties at fixed halo mass. We find this a relevant and interesting question to study, as recent works show that halo assembly bias can affect the selection of galaxies in modern surveys [see e.g., [81]]. Seeing how strong of a dependence on halo assembly bias we find in AbacusPNG and AbacusSummit would give us insight into the importance of developing models for b_ϕ that take into account dependencies beyond linear bias (b_1), or interchangeably, halo mass (M_h).

A. Assembly bias properties

We first start by introducing some of the most ubiquitously studied halo assembly bias properties: concentration, shear (environment), and accretion rate.

1. Concentration

The link between halo concentration and accretion history has been studied extensively in the literature [82–85]. It has

been shown that recent merger activity induces dramatic changes in halo concentrations, and that these responses linger over a period of several dynamical times, corresponding to many Gyr [see, e.g., [86]]. Relevant to assembly bias studies is the fact that halo concentration has a bearing on both the halo occupation distribution and the halo clustering [e.g., [81,84,87–89]].

In this work, we adopt the following proxy for the concentration of each halo:

$$c = r_{90}/r_{25}, \qquad (27)$$

following the recommendation of Hadzhiyska *et al.* [90], where r_{90} and r_{25} are defined as the radii, within which 90% and 25% of the halo particles are contained inside a sphere centered on the halo center. For more details on the halo finder and virial mass definition, we refer the reader to Bryan and Norman [91] and Hadzhiyska *et al.* [90].

2. Shear

Our procedure for obtaining the adaptive halo shear ("shear," for short) manipulates the smoothed particle density field into the shear field [see e.g., [92]]. Namely, to calculate the local "shear" around a halo, we first compute a dimensionless version of the tidal tensor, defined as:

$$T_{ij} \equiv \partial^2 \phi_R / \partial x_i \partial x_j, \qquad (28)$$

where ϕ_R is the dimensionless potential field calculated using Poisson's equation: $\nabla^2 \phi_R = -\rho_R/\bar{\rho}$ (the subscript R

corresponds to the choice of smoothing scale). We then calculate the tidal shear q_R^2 via:

$$q_R^2 \equiv \frac{1}{2}[(\lambda_2 - \lambda_1)^2 + (\lambda_3 - \lambda_1)^2 + (\lambda_3 - \lambda_2)^2], \qquad (29)$$

where λ_i are the eigenvalues of T_{ij}. Physically, the shear at some particular point in space measures the amount of anisotropic pulling due to gravity at a given point in space.

3. Accretion rate

Previous studies of local-type PNG have surmised that the b_ϕ parameter is sensitive to the accretion rate of a halo, as the coupling of long- and short-wavelength perturbations is sensitive to the formation epoch of halos/galaxies. In this study, we define the accretion rate as follows:

$$\Gamma_{\mathrm{dyn}}(t) = \frac{\Delta \log(M)}{\Delta \log(a)} = \frac{\log[M(t)] - \log[M(t - t_{\mathrm{dyn}})]}{\log[a(t)] - \log[a(t - t_{\mathrm{dyn}})]}, \qquad (30)$$

where $M \equiv M_{200\,\mathrm{m}}$ (mass within 200 times the mean density of the Universe) and the subscript "dyn" refers to the dynamical time, defined as:

$$t_{\mathrm{dyn}}(z) \equiv t_{\mathrm{cross}}(z) = \frac{2R_\Delta}{v_\Delta} = 2^{3/2} t_H(z) \left(\frac{\rho_\Delta(z)}{\rho_c(z)}\right)^{-1/2}, \qquad (31)$$

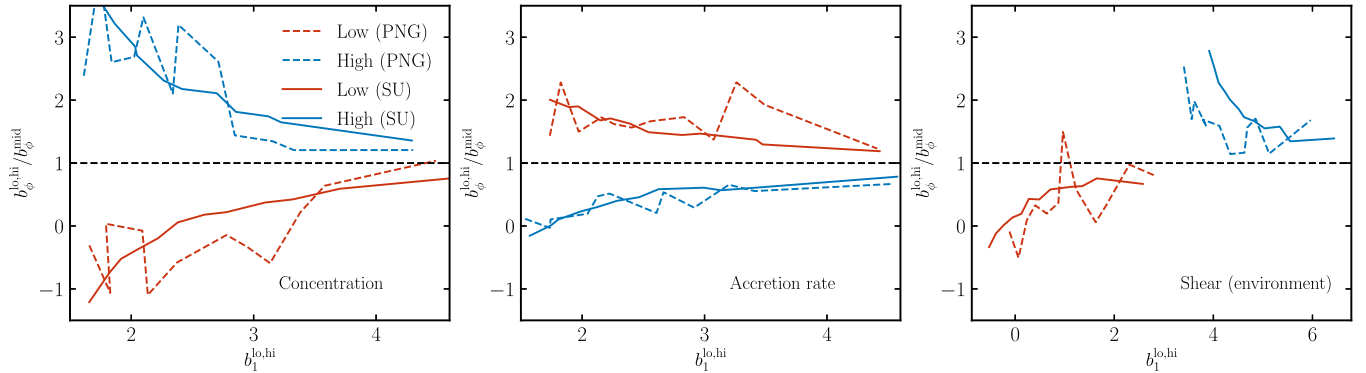where $t_H(z) \equiv 1/H(z)$ is the Hubble time.

FIG. 4. Response of the primordial bias parameter b_ϕ to the halo assembly bias properties: concentration [see Eq. (27)], accretion rate [see Eq. (30)] and shear [see Eq. (29)]. The dashed curves come from the AbacusPNG set of simulations, whereas the solid curves come from the separate universe pair of simulations, which uses the original AbacusSummit suite. The halos are split into 12 mass bins and then each bin is further split into 3 bins of equal sizes (33%, 66% percentile) based on the secondary property being considered (low, mid, high). Here, for all mass bins, we report the pairs $[b_1^{\mathrm{high}}, b_\phi^{\mathrm{high}}/b_\phi^{\mathrm{mid}}]$ in blue and $[b_1^{\mathrm{low}}, b_\phi^{\mathrm{low}}/b_\phi^{\mathrm{mid}}]$ in red. The separate universe and AbacusPNG curves are in good agreement with each other. For all three properties, the low-mass (low-bias) bins exhibit the largest variations in their values of b_ϕ. If a given galaxy sample preferentially occupies low- or high-concentration (accretion rate) halos, then that may not reflect on the inference from the two-halo clustering, but it will affect PNG analysis. The concentration response is stronger than that of the accretion rate, and that the trend is reversed between the two, which makes sense, as actively accreting objects have lower concentration (have more spread out substructure). We note that a different definition of accretion rate might yield a stronger response. To see a change in b_ϕ for shear, one needs to vary the linear bias significantly at fixed halo mass.

B. Response to PNG

In Fig. 4, we show the response of the primordial bias parameter b_ϕ to the halo assembly bias properties: concentration [see Eq. (27)], accretion rate [see Eq. (30)] and shear [see Eq. (29)]. The dashed curves come from the AbacusPNG set of simulations, whereas the solid curves come from the separate universe pair of simulations, which uses the original AbacusSummit suite. The halos are split into 12 mass bins and then each bin is further split into 3 bins of equal sizes (33%, 66% percentile) based on the secondary property being considered (low, mid, high). Here, for all mass bins, we report the pairs $[b_1^{\rm high}, b_\phi^{\rm high}/b_\phi^{\rm mid}]$ in blue and $[b_1^{\rm low}, b_\phi^{\rm low}/b_\phi^{\rm mid}]$ in red. We note several interesting features. To convert halo mass into linear bias, we adopt the empirical relation from Tinker *et al.* [79] and calculate the linear bias of each halo before computing the statistics of interest.

The separate universe and AbacusPNG curves are in good agreement with each other, though as noted before (see Fig. 3), the fitting method yields much noisier results. We note that some of the differences between the two may be due to resolution effects (which might impact the lowest- and highest-halo masses).

For all three properties, the low-mass (low-bias) bins exhibit the largest variations in their values of b_ϕ. In the case of concentration and accretion rate, the response of b_ϕ to the halo property is evidently much stronger compared with the response of b_1. This is crucial for modeling scale-dependent bias, as it suggests that if a given galaxy sample preferentially occupies low- or high-concentration (accretion rate) halos, then that may not reflect on the inference from the two-halo clustering (which is largely insensitive to these properties), but it will affect PNG analysis. This indicates we should aim to perform careful small-scale analysis in order to constrain the galaxy-halo connection.

We further notice that the concentration response is stronger than that of the accretion rate, and that the trend is reversed between the two, which makes sense, as actively accreting objects have lower concentration (have more spread out substructure). On the other hand, for the case of shear, for which high-shear corresponds to high b_ϕ, we find very different behavior: to see a change in b_ϕ, one needs to vary the linear bias significantly at fixed halo mass. This is possible if we find that environment is indeed an important factor to consider for improving mass-only galaxy population models and resolving lensing-clustering tensions such as "lensing is low" [e.g., [93–95]].

V. DESI GALAXIES

In this section, we explore realistic samples of DESI-like galaxies for the two tracers most relevant for PNG: luminous red galaxies (LRGs) and quasistellar objects (QSOs). Specifically, using halo occupation distribution (HOD) fits to early data from DESI, we investigate how b_ϕ

responds to various extensions of the standard HOD parametrization, which are allowed by the data. We then propagate the uncertainty of measuring b_ϕ into an uncertainty on $f_{\rm NL}$ by making simple Fisher forecasts on the combination of $b_\phi f_{\rm NL}$ and report the marginalized constraints on $f_{\rm NL}$ for the two tracers.

A. HOD of DESI galaxies

Here, we summarize the HOD model used to fit the LRG and QSO clustering with DESI 1% data [see [25]]. These are the parameter values we adopt when creating our synthetic LRG and QSO catalogs. We note that we apply this HOD model (with and without extensions) to both the separate universe as well as the AbacusPNG simulations in order to estimate b_ϕ and derive constraints on the combination $b_\phi f_{\rm NL}$.

DESI targets LRGs at $z \lesssim 1$, as they are bright galaxies with a prominent break at 4000 Å in their spectra, which allows them to be selected relatively easily in the data. In addition, they are highly biased tracers, which makes their BAO feature more prominent compared with other galaxy types. The DESI LRG sample has a fairly constant number density between $0.4 < z < 0.8$ of approximately 5×10^{-4} [Mpc/h]$^{-3}$. QSOs (quasars), on the other hand, are the tracer choice for studying large-scale structures at high redshifts due to their extremely high luminosities. Their number density is roughly constant between $0.8 < z < 2.1$, at 2×10^{-5} [Mpc/h]$^{-3}$.

The vanilla HOD model for the two samples is given by the standard formalism of Zheng *et al.* [96]:

$$\bar{n}_{\rm cent}(M) = \frac{f_{\rm ic}}{2} \, {\rm erfc} \left[\frac{\log_{10}(M_{\rm cut}/M)}{\sqrt{2}\sigma} \right], \qquad (32)$$

$$\bar{n}_{\rm sat}(M) = \left[\frac{M - \kappa M_{\rm cut}}{M_1} \right]^\alpha \qquad (33)$$

where we note that the LRG satellite occupations equation has a modification in the form of multiplication by $\bar{n}_{\rm cent}(M)$. $M_{\rm cut}$ determines the minimum mass of a halo to host a central galaxy, M_1 sets the pivot scale of the power law of satellite occupation, σ controls the steepness of the transition from 0 to 1 in the number of central galaxies, α is the power law index on the number of satellite galaxies, $\kappa M_{\rm cut}$ gives the minimum halo mass to host a satellite galaxy, $f_{\rm ic}$, which is a downsampling factor controlling the overall number density of the mock galaxies.

In the vanilla model, the velocity of the central galaxy is taken as the average velocity of the so-called "L2" subhalo [see [90]]. For the satellite galaxies, the velocities are inherited from random halo particles. The analysis of Yuan *et al.* [25] also includes velocity bias, which is necessary for modeling redshift-space clustering on small scales [e.g. [22,97]]. In the AbacusHOD model, it is parametrized as:

TABLE III. Marginalized posteriors of the HOD parameters from fits to measurements of the clustering, $\xi(r_p, r_\pi)$, of DESI LRGs and QSOs in the DESI 1% survey, performed in Yuan *et al.* [25]. For our fiducial galaxy catalogs, we adopt the best fit values of these fits at redshifts $z = 0.5$, $z = 0.8$, and $z = 1.4$ for the three samples shown here, adopting the same vanilla HOD model of AbacusHOD, and the same set of simulations, which ensures consistency of the halo mass and HOD parameter definitions. In addition, we explore extensions to the vanilla HOD model in the form of concentration and environment dependence. The error bars show 1σ uncertainties. Units of mass are given in $h^{-1}M_\odot$.

Tracer	LRG	LRG	QSO
Redshift	$0.4 < z < 0.6$	$0.6 < z < 0.8$	$0.8 < z < 2.1$
$\log M_{\text{cut}}$	$12.79^{+0.15}_{-0.07}$	$12.64^{+0.17}_{-0.05}$	$12.2^{+0.6}_{-0.1}$
$\log M_1$	$13.88^{+0.11}_{-0.11}$	$13.71^{+0.07}_{-0.07}$	$14.7^{+0.6}_{-0.6}$
σ	$0.21^{+0.11}_{-0.10}$	$0.09^{+0.09}_{-0.05}$	$0.12^{+0.28}_{-0.06}$
α	$1.07^{+0.13}_{-0.16}$	$1.18^{+0.08}_{-0.13}$	$0.8^{+0.4}_{-0.2}$
κ	$1.4^{+0.6}_{-0.5}$	$0.6^{+0.4}_{-0.2}$	$0.6^{+0.8}_{-0.2}$
α_c	$0.33^{+0.05}_{-0.07}$	$0.19^{+0.06}_{-0.09}$	$1.54^{+0.17}_{-0.08}$
α_s	$0.80^{+0.07}_{-0.07}$	$0.95^{+0.07}_{-0.06}$	$0.6^{+0.6}_{-0.3}$

(i) $\alpha_{\text{vel,c}}$ modulates the peculiar velocity of the central galaxy relative to the halo center along the line-of-sight (LOS):

$$v_{\text{cent,z}} = v_{\text{L2,z}} + \alpha_{\text{vel,c}} \delta v(\sigma_{\text{LoS}}), \qquad (34)$$

where $v_{\text{L2,z}}$ denotes the LOS component of the central subhalo velocity, $\delta v(\sigma_{\text{LoS}})$ denotes the Gaussian scatter, and $\alpha_{\text{vel,c}}$ is the central velocity bias parameter.

(ii) $\alpha_{\text{vel,s}}$ modulates how the satellite galaxy peculiar velocity deviates from that of its host particle:

$$v_{\text{sat,z}} = v_{\text{L2,z}} + \alpha_{\text{vel,s}}(v_{\text{p,z}} - v_{\text{L2,z}}), \qquad (35)$$

where $v_{\text{p,z}}$ denotes the line-of-sight component of particle velocity, and $\alpha_{\text{vel,s}}$ is the satellite velocity bias parameter.

In Table III, we show the marginalized posteriors of the HOD parameters from fits to measurements of the clustering, $\xi(r_p, r_\pi)$, of DESI LRGs and QSOs in the DESI 1% survey, performed in Yuan *et al.* [25]. The error bars show 1σ uncertainties. Units of mass are given in $h^{-1}M_\odot$. For our fiducial galaxy catalogs, we adopt the best fit values of these fits at redshifts $z = 0.5$, $z = 0.8$, and $z = 1.4$ for the three samples shown here, using the AbacusSummit boxes to ensure consistency of the halo mass and HOD parameter definitions.

B. Galaxy assembly bias extensions to the model

In this analysis, we study the effect of allowing physically motivated extensions on the inferred b_ϕ from the

separate universe approach. This section summarizes the assembly bias extensions allowed in the AbacusHOD model (see [22] for more details):

(i) A_{cent} or A_{sat} are the concentration-based secondary bias parameters for centrals and satellites, respectively. $A_{\text{cent,sat}} = 0$ indicate no concentration-based secondary bias. Positive values of A indicate a preference for lower concentration halos, and vice versa, at fixed halo mass. The concentration definition adopted here is equivalent to the one in Sec. IV.

(ii) B_{cent} or B_{sat} are the environment-based secondary bias parameters for centrals and satellites, respectively. The environment is defined as the mass density within a $r_{\text{env}} = 5$ Mpc/h top hat of the halo center, excluding the halo itself, which we note differs from the "shear" parameter we define in Sec. IV A. Nonetheless, we check that qualitatively these two environment parameters yield very similar results, and we choose this parametrization for consistency with the DESI analysis. $B_{\text{cent,sat}} = 0$ indicate no environment-based secondary bias. Positive values of B indicate a preference for halos in less dense environments, and vice versa, at fixed halo mass.

In this study, we are interested in exploring how the parameter b_ϕ changes, as we include assembly bias extensions to the vanilla HOD model. In particular, we follow the procedure outlined below:

We first adopt the best-fit values from the $\xi(r_p, r_\pi)$ to the DESI 1% data (see Table III), we construct our fiducial samples of LRGs at $z = 0.5$ and $z = 0.8$, and QSOs at $z = 1.4$. We refer to this as the fiducial or best-fit sample from hereon.

We then create mock catalogs in a grid centered on the fiducial parameter values. In particular, at each grid point, we vary the two mass parameters $\log M_{\text{cut}}$ and $\log M_1$ by ± 0.5, which roughly control the amplitude of the 2-halo and the transitioning between the 1- and 2-halo regimes, as well as the central galaxy concentration and assembly bias parameters A_{cent} and B_{cent} by ± 0.3. These ranges correspond to the prior choices on these parameters in Yuan *et al.* [25]. For each parameter, we explore 5 linearly spaced values in the specified ranges (corresponding to a step size of 0.25 in the mass parameters and 0.15 in the assembly bias parameters).

We use mock covariance matrices generated from the 1800 small AbacusSummit boxes for our three samples [see [25]] from $\sim$2000 AbacusSummit small boxes by measuring the redshift-space 2D correlation function, $\xi(r_p, r_\pi)$, which can be obtained via the Landy and Szalay [98] estimator:

$$\xi(r_p, r_\pi) = \frac{DD - 2DR + RR}{RR}, \qquad (36)$$

where r_p and r_π are transverse and LOS separations in comoving units, respectively, whereas DD, DR, and RR are

the normalized data-data, data-random, and random-random pair counts in each bin of (r_p, r_π). In principle, $\xi(r_p, r_\pi)$ captures the full information content of the two-point clustering and can thus yield stronger constraints on the galaxy-halo connection [22] compared with projected statistics such as $w_p(r_p)$, which wash out the LOS information. Here, we measure $\xi(r_p, r_\pi)$ in 14 logarithmically spaced bins between $\log r_{p,\min} = -0.83$ and $\log r_{p,\max} = 1.5$ in the transverse direction and $r_{\pi,\max} = 32$ Mpc/h with $\Delta r_\pi = 4$ Mpc/h in the LOS direction. Thus, provided no other cuts are made (e.g., to ensure nonsingularity of the covariance matrix), the number of degrees of freedom (d.o.f.) is $\sim$112 (note that we lose 4 d.o.f. due to the parameters being varied). We rescale the small-box $(0.5^3 \, [\text{Gpc}/h]^3)$ covariance matrices by a factor of 20, 30, and 200 for the three samples, respectively, to roughly reflect the volume of DESI Y5 for these tracers. We estimate that the intrinsic uncertainty of the HOD model in constraining the galaxy-halo connection is roughly compensated by the effective volume ratio of DESI Y5, and leave the discussion of HOD model uncertainties for future work.

We compute $\xi(r_p, r_\pi)$ for the $5^4 = 625$ grid points and using the computed mock covariance matrices, calculate the "distance" of each sample from the "true" $\xi^{\text{fid}}(r_p, r_\pi)$ corresponding to the fiducial HOD sample defined in Table III:

$$\Delta\chi^2 = (\boldsymbol{\xi} - \boldsymbol{\xi}^{\text{fid}})^T \mathcal{C}^{-1} (\boldsymbol{\xi} - \boldsymbol{\xi}^{\text{fid}}). \tag{37}$$

Finally, we select all samples from the 625-point grid, which have $\Delta\chi^2 \leq$ d.o.f., which corresponds to roughly 1σ deviations from the fiducial sample clustering (which is equivalent to the best-fit sample), and calculate b_ϕ for each, adopting the separate universe approach. In Fig. 5, we show a scatter plot of b_ϕ and mean halo mass for all three samples for the points that pass the $\Delta\chi^2 \leq$ d.o.f. criterion. The mean and fiducial values are shown as a cross and larger circle, respectively. We also show the universality relation with $c = 0.8$ [see Eq. (17)] for each sample. Intriguingly, the galaxy samples seem to prefer a lower value of c than in the case of the halos (cf. $c = 0.9$). We see that the samples that satisfy the condition of $\Delta\chi^2 \leq$ d.o.f. (shown as dots on the plot) follow a thin slanted contour on the b_ϕ–$\bar{M}_{\text{halo}}$ plane. We attribute this to the fact that the mean halo mass (which is a proxy of linear bias) is relatively well constrained compared with b_ϕ. A closer look at the surviving HOD samples tells us that most of the scatter in b_ϕ is due to variations in the concentration parameter, A_{cent}. Indeed, as seen in Fig. 4, concentration has a stronger impact on b_ϕ than on the linear bias. It is also rather poorly constrained compared with the mass parameters due to its subtle effect on the clustering. The fiducial, mean, half difference between the maximum and minimum, and standard deviation computed using the eligible values of the
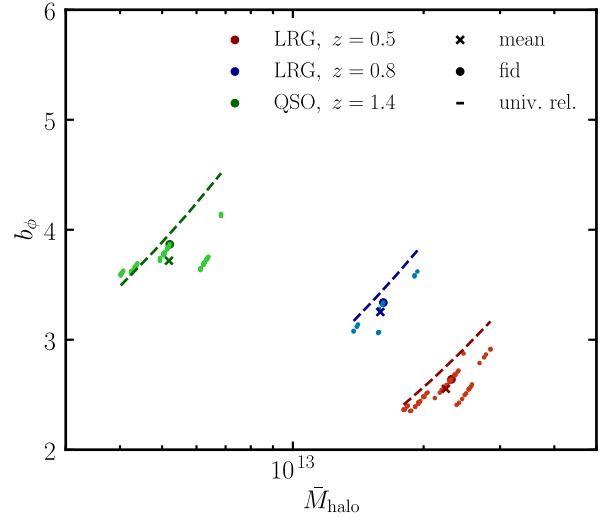

FIG. 5. Scatter plot of b_ϕ and mean halo mass for the three samples of interest to this study: LRGs at $z = 0.5$ (red) and $z = 0.8$ (blue), and QSOs at $z = 1.4$ (green). Dots are shown for the HOD samples that satisfy the condition $\Delta\chi^2 \leq$ d.o.f. as described in Sec. V B. We show the mean and fiducial values for each sample with a cross and large circle, respectively, as well as the universality relation (dashed) with $c = 0.8$ [see Eq. (17); note that a lower value of c is found for the galaxies compared with the halos, $c = 0.9$]. We see that the samples follow a thin slanted contour on the b_ϕ–$\bar{M}_{\text{halo}}$ plane. We attribute this to the fact that the mean halo mass (which is a proxy of linear bias) is relatively well constrained compared with b_ϕ. The standard deviation of b_ϕ for the three samples is roughly 11%, 6% and 5%, respectively.

parameter b_ϕ (i.e., from the samples that are within 1σ of the best-fit) for the LRGs at $z = 0.5$, $z = 0.8$, and QSOs at $z = 1.4$ is

$$\{b_\phi^{\text{fid}}, b_\phi^{\text{mean}}, \Delta b_\phi/2, b_\phi^{\text{std}}\}$$

$$= \{2.64, 2.55, 0.29, 0.13\} \quad \text{(for LRG}, z = 0.5)$$

$$\{3.34, 3.25, 0.19, 0.11\} \quad \text{(for LRG}, z = 0.8)$$

$$\{3.87, 3.72, 0.19, 0.10\} \quad \text{(for QSO}, z = 1.4) \tag{38}$$

This corresponds to an uncertainty of 11%, 6%, and 5%, respectively. We additionally note that if we were able to pinpoint the mass of the tracers, the uncertainty on b_ϕ would decrease noticeably. While it is difficult to disentangle assembly bias and halo mass from clustering data alone, including weak lensing constraints on the galaxy-halo link would substantially help us to break HOD parameter degeneracies, allowing us to put much tighter constraints on b_ϕ.

Note: We expect most of the assembly bias signal to be contained in the central galaxy population [see e.g., [92]]. We leave the full analysis that includes running chains with all HOD extensions for later work. Beyond that, there may be other halo properties that are correlated with b_ϕ, which

may need to be further studied. It is also possible that the satellite population, which makes up between 10% and 20%, of all galaxies of interest to this study is sensitive to halo properties, which are degenerate with b_ϕ, such as halo concentration and accretion rate. We note that future data releases of the clustering of DESI LRGs and QSOs will feature much smaller error bars, which will also improve our constraints on the galaxy-halo connection and on the assembly bias of these samples.

C. Fisher formalism

Here, we provide a short revision of the Fisher formalism [99]. For a given data vector, $\mathcal{D}$, and a vector of parameters, $\boldsymbol{\theta} = \{\theta_\alpha\}$ with α denoting the index of each parameter, we can define the Fisher matrix as:

$$F_{\alpha\beta} = \left\langle \frac{\partial \log \mathcal{L}(\mathbf{D}|\boldsymbol{\theta})}{\partial \theta_\alpha} \frac{\partial \log \mathcal{L}(\mathbf{D}|\boldsymbol{\theta})}{\partial \theta_\beta} \right\rangle, \qquad (39)$$

where $\mathcal{L}(D|\boldsymbol{\theta})$ is the likelihood. This Fisher matrix can then be used to obtain a minimum estimate of the error on each of our parameters θ_α [100]:

$$\sigma[\theta_\alpha] = \sqrt{(F^{-1})_{\alpha\alpha}}, \qquad (40)$$

where $\sigma[\cdot]$ defines the error on our parameter of interest. Thus, the inverse of the Fisher information lets us infer the maximum information that we can obtain about each of our theory parameters, given some observed data. We note that $\sigma[\theta_\alpha] = F_{\alpha\alpha}^{-1/2}$ yields the nonmarginalized error on the parameter θ_α.

In this work, we are interested in the halo and galaxy power spectrum, which we calculate using the ABACUSUTILS package [101]. Assuming that their likelihoods are Gaussian distributed, which works particularly well for the power spectrum on the large scales we are interested in,[6] we can express the log-likelihood as:

$$\log \mathcal{L}(\mathbf{D}|\boldsymbol{\theta}) = -\frac{1}{2} \sum_{ij} (P(k_i) - \bar{P}(k_i))\mathcal{C}_{ij}^{-1}(P(k_j) - \bar{P}(k_j)),$$

$$(41)$$

where i and j sum over all measurement bins, P denotes the theory power spectrum, $\bar{P}$ is the observed power spectrum, and $\mathcal{C}$ is the covariance matrix.

Assuming that the covariance matrix is independent of the parameters, θ_α, we can further simplify the Fisher matrix expression as:

$$F_{\alpha\beta} = \frac{\partial \bar{P}(k_i)}{\partial \theta_\alpha} \mathcal{C}_{ij}^{-1} \frac{\partial \bar{P}(k_j)}{\partial \theta_\beta}. \qquad (42)$$

---

[6] We also note that to get an estimate of the information content, this approximation is good enough.

Thus, we only require the derivative of our summary statistic with respect to the parameters and its covariance to assess its information content. We note that in our case the covariance matrix does depend on the cosmological parameters, but neglecting this dependence gives a better approximation for the true information content Carron [102]. The details of how the covariance matrix is computed in this study are presented in Sec. V B.

D. Constraints on $b_\phi f_{\mathrm{NL}}$

In this study, we are interested in assessing the information content of the scale-dependent bias and constraining f_{NL} by robustly marginalizing over b_ϕ.

To compute the Fisher matrix, we need a covariance matrix of the power spectrum. Here, we use an analytically calculated autopower spectrum covariance matrices [103] specifically tailored to the DESI LRG and QSO tracers at $z = 0.4$–0.6 and $z = 1.1$–1.4, respectively. We obtain the cross-power spectrum covariance matrix by adopting the Gaussian approximation, which should hold on large scales, and rescaling the matrix to match the Y5 volume for each tracer in each redshift range. As validation, we compute the derivatives in two different ways: numerically (by taking finite differences) and analytically (by assuming linear theory) and find that they are in excellent agreement on large scales. We can obtain the numerical derivative with respect to $b_\phi f_{\mathrm{NL}}$ as:

$$\frac{\partial \bar{P}(k_i)}{\partial [b_\phi f_{\mathrm{NL}}]} \approx \frac{\bar{P}(k_i)|_{[b_\phi f_{\mathrm{NL}}]=100 b_\phi^*} - \bar{P}(k_i)|_{[b_\phi f_{\mathrm{NL}}]=-100 b_\phi^*}}{2\delta[b_\phi f_{\mathrm{NL}}]}, \quad (43)$$

where $\delta[b_\phi f_{\mathrm{NL}}] = 100 b_\phi^*$ for the simulations used in this section, AbacusPNG_c30{2,3}_ph000. The values of b_ϕ^* we adopt for the three samples (LRGs at $z = 0.5$, $z = 0.8$, and QSOs at $z = 1.4$) are as follows: 2.64, 3.34, 3.87. We compute the Fisher matrix both with and without redshift space distortions, but find that including this effect has negligible impact on the f_{NL} constraints, as most of the information comes from large scales, which are unaffected by redshift space distortions. In redshift space, we employ the Legendre multipole expansion of the redshift-distorted power spectrum, $P(k, \mu)$:

$$P_\ell(k) = \frac{2\ell + 1}{2} \int_{-1}^{1} P(k, \mu) L_\ell(\mu) d/\mu \qquad (44)$$

where L_ℓ is the Legendre polynomial and $P_\ell(k)$ are the multipoles of the redshift-distorted power spectrum, $P(k, \mu)$. From hereon, we drop the ℓ subscript to simplify the notation.

We also perform direct fits to the simulation measurements. For the autopower spectrum, we vary the combination $b_\phi f_{\mathrm{NL}}$, the linear bias b_1, and the shot noise parameter a (to allow for deviations from the assumption

of Poisson shot noise), whereas for the cross-power spectrum case, we vary only $b_\phi f_{NL}$ and b_1.

$$P_{gg}(k) = \left[b_1 + \frac{b_\phi f_{NL}}{\mathcal{M}(k,z)}\right]^2 P_{mm}(k) + \frac{1}{\bar{n}_g}(1+a)$$

$$P_{gm}(k) = \left[b_1 + \frac{b_\phi f_{NL}}{\mathcal{M}(k,z)}\right]P_{mm}(k), \tag{45}$$

where $\bar{n}_g$ is the mean number density of the galaxy sample and $\mathcal{M}(k,z)$ is defined in Eq. (5). The linear theory approximation adopted here starts to break down beyond $k \gtrsim 0.15h/\mathrm{Mpc}$. Indeed, in that regime, we find that the inferred values of $b_\phi f_{NL}$ start to substantially deviate from the simulation "truth." We note that while the cross-power spectrum, $P_{gm}(k)$, is not directly observable, we can access the galaxy-matter projected clustering via joint probes with weak lensing surveys and the CMB map.

We present our constraints on $b_\phi f_{NL}$ in Fig. 6. The solid line is obtained via a direct fit (using `curvefit`) to the linear theory model in redshift space, $P_{\ell=0,2,4}$, whereas the dashed line is obtained via the Fisher approximation (nonmarginalized). We see that for the gg case, the agreement between Fisher and the direct fit is very good, though for small values of k_{max}, there is a larger gap between the two. We attribute this to the larger noise in the fit, as there are much fewer modes available, and the PNG component is degenerate with the linear bias. For the gm case, the agreement is poorer, which we attribute to the approximation adopted for the cross-power spectrum covariance matrix. As we increase k_{max}, the constraints on $b_\phi f_{NL}$ improve only marginally, since the small-scale power spectrum is negligibly affected by the presence of local-type PNG. The constraints for the three tracers from the autopower spectra are

$$\sigma[b_\phi f_{NL}] = 25, 20, 18, \tag{46}$$

respectively.

E. Constraints on f_{NL}

Taking into account our findings from Sec. V B on the uncertainty of b_ϕ, we can now translate the uncertainty on the combination $b_\phi f_{NL}$ into an uncertainty on f_{NL} using linear error propagation, as follows:

$$\sigma[f_{NL}] = f_{NL}\left(\left(\frac{\sigma[b_\phi]}{b_\phi}\right)^2 + \left(\frac{\sigma[b_\phi f_{NL}]}{b_\phi f_{NL}}\right)^2 \right.$$
$$\left. - 2\frac{\sigma[b_\phi, f_{NL}]}{b_\phi f_{NL}}\right)^{1/2}. \tag{47}$$

Assuming that b_ϕ and f_{NL} have negligible covariance, we can simplify the equation above by ignoring the cross-term, $\sigma[b_\phi, f_{NL}] \approx 0$. Substituting the standard deviation
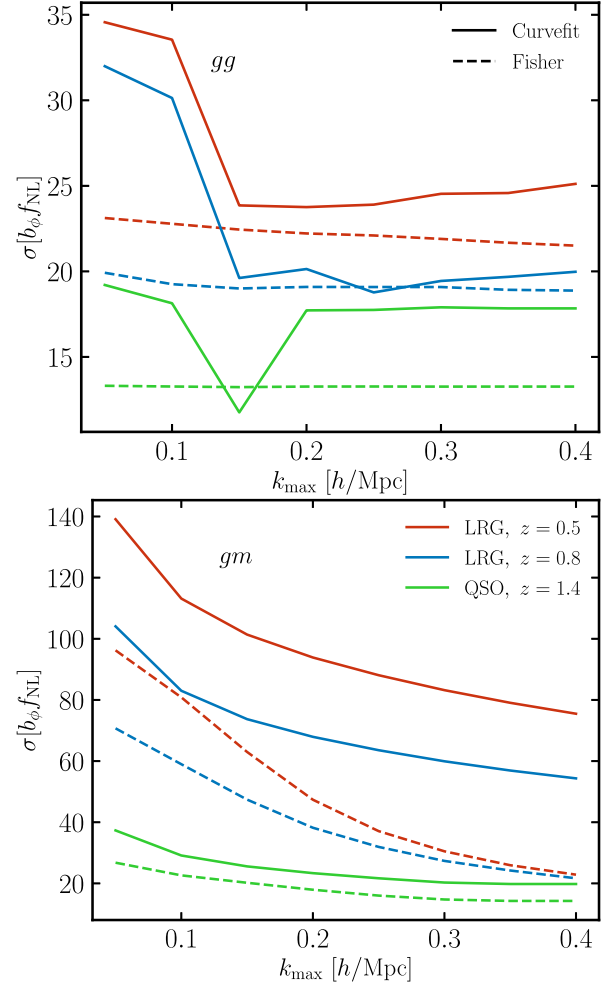

FIG. 6. Constraints on the combination $b_\phi f_{NL}$ as a function of scale (k_{max}) for the galaxy auto (gg, upper panel) and cross- (gm, lower panel) power spectrum for the LRGs at $z = 0.5$ (red) and $z = 0.8$ (blue) and QSOs at $z = 1.4$ (green). The solid line is obtained via a direct fit (`curvefit`) to the linear theory model in redshift space, $P_{\ell=0,2,4}$, whereas the dashed line is obtained via the Fisher approximation (nonmarginalized). We see that for the gg case, the agreement between Fisher and the direct fit is very good, though for small values of k_{max}, for which there are very few modes available for the fit, there is a larger gap between the two. For the gm case, the agreement is poorer, which we attribute to the approximation adopted for the cross-power spectrum covariance matrix. As we increase k_{max}, the constraints on $b_\phi f_{NL}$ improve only marginally, since the small-scale power spectrum is not affected by the presence of local-type PNG. The constraints for the three tracers from the autopower spectra are $\sigma[b_\phi f_{NL}] = 25, 20, 18$, respectively.

and fiducial values of b_ϕ for the three samples from Sec. V B and the $\sigma[b_\phi f_{NL}]$ from the Fisher analysis, we obtain the following constraints on f_{NL} from the galaxy autopower spectrum:

$$\sigma[f_{NL}] = 15, 8, 7. \tag{48}$$

In this analysis, we assume that the other cosmological parameters are kept fixed. These constraints are shown in Fig. 7, where we see the constraints as a function of scale (k_{max}). The best constraints are obtained from the QSOs at $z = 1.4$, as expected, as quasars have the largest volume and
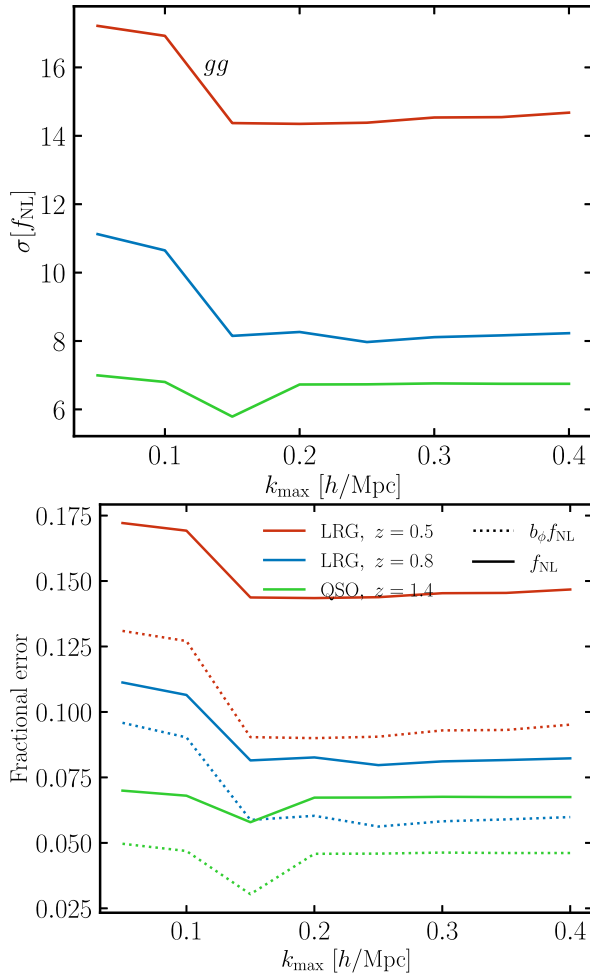

FIG. 7.   Upper panel: constraints on f_{NL} as a function of scale (k_{max}) for the galaxy autopower spectrum of the LRGs at $z = 0.5$ (red) and $z = 0.8$ (blue) and QSOs at $z = 1.4$ (green). We see that the best constraints are obtained from the LRGs at $z = 0.8$, which have a larger volume. While the QSOs had poorer constraints on $b_\phi f_{NL}$ compared with the other two tracers, we see that they yield comparable $\sigma[f_{NL}]$ to the LRGs at $z = 0.5$, as b_ϕ is highest for that sample. We quote the minimum error on f_{NL} after taking into account the uncertainty of b_ϕ for each of the three samples as $\sigma[f_{NL}] = 15, 8, 7$, respectively. Lower panel: Fractional error, defined as $\sigma[\cdot]/[\cdot]$, of the combination $b_\phi f_{NL}$ (dotted line) and just f_{NL} using Eq. (47) and the constraints on b_ϕ from Sec. V B. For the LRGs at $z = 0.8$ and the QSOs, the fractional error on f_{NL} compared with that on $b_\phi f_{NL}$ worsens by about a factor of 1.5 due to the large uncertainty on the astrophysical parameter, b_ϕ, and by 2 for the LRGs at $z = 0.5$. This shows that it is of great importance to understand the galaxy-halo connection of the DESI samples in order to reduce both the uncertainty and the bias on the inferred value of f_{NL}.

largest value of b_ϕ. The LRGs at $z = 0.8$ yield very comparable constraints, as their b_ϕ is relatively high and well constrained. In the figure, we also display the fractional error, $\sigma[\cdot]/[\cdot]$, of the combination $b_\phi f_{NL}$ and just f_{NL} using Eq. (47) and the constraints on b_ϕ from Sec. V B. For the LRGs at $z = 0.8$ and the QSOs, the fractional error on f_{NL} compared with that on $b_\phi f_{NL}$ is a factor of 1.5 worse due to the large uncertainty on the astrophysical parameter, b_ϕ. For the LRGs at $z = 0.5$, the two curves differ by almost a factor of 2, as b_ϕ is poorly constrained for that sample (see Fig. 5). This finding hints that more investigation into the complex link between galaxy formation and local-type PNG is warranted. In order to place $\sigma[f_{NL}] \sim 1$ constraints on this model of the primordial Universe, we need to not only make precise measurements on large-scales using spectroscopic or weak lensing surveys of large volumes, but also substantially improve our constraints on the galaxy-halo connection and thus, b_ϕ.

VI. SUMMARY AND CONCLUSIONS

Observations of the CMB have revolutionized our understanding of the primordial Universe, demonstrating that primordial fluctuations were nearly homogeneous and isotropic with an almost scale-invariant power spectrum. To explain these observations, a number of early Universe theories such as inflationary and ekpyrotic models have been proposed. In order to discriminate between these models, we need to attain definitive evidence of their predicted byproducts such as signatures of primordial gravitational waves in the B-mode polarization of the CMB and primordial non-Gaussianity (PNG) [42,104,105]. PNG is an invaluable probe, as it can reveal information about the field content, dynamics, and strength of interactions in the early Universe. While measurements of the CMB have placed tight constraints on the presence of PNG, large-scale structure probes are potentially even more powerful, as they can access a larger number of Fourier modes corresponding to the volume covered by the survey.

In the next few years, large-scale structure experiments such as DESI and *Roman* will catalogue the 3D positions of tens of millions of galaxies, mapping an unprecedentedly large volume of space. Extracting subpercent comparisons between survey observations and cosmological predictions of PNG requires high-precision mock data and robust theoretical models. Although approximate methods are capable of generating sample data, cosmological N-body simulations are a central tool for making accurate forecasts for the nonlinear regime of gravitational structure formation. Moreover, extracting PNG constraints from large-scale structure is challenging due to complex observational effects (such as survey depth, mask, fiber collisions), modeling of the summary statistics, and disentangling astrophysical from primordial physics effects. Understanding the relation between PNG and late-time

observables is the main objective of the AbacusPNG set of N-body simulations, presented in this work.

Here, we present a new set of simulations, AbacusPNG, which was run with the extremely fast and accurate Abacus code [24], modified to possess a non-Gaussian primordial gravitational potential. The aim of this work is to aid current and next generation of surveys in constraining local-type PNG from large-scale structure and to validate the simulation methods for future larger simulation suites with Abacus. First, we have introduced the AbacusPNG set in Sec. II. Next, we have validated its initial conditions by studying the matter bispectrum and comparing it against the tree-level prediction (see Fig. 1). In addition, we have compared the matter power spectrum at late times with the one-loop theoretical prediction, finding good agreement on the relevant scales (see Fig. 2). Finally, we have tested that the b_ϕ we compute from the PNG simulations is in agreement with the value of the parameter inferred using the separate universe technique on the original linear-derivative AbacusSummit boxes (see Fig. 3). We also quote, the slope of the universality relation for the Abacus catalogs as $c \approx 0.85$ [see Eq. (17)].

As a science case, we focus on the halo and galaxy fields in Secs. IV and V, respectively. First, we study the halo response to different assembly bias properties, which have been associated with the galaxy samples of modern surveys such as BOSS and DESI [see [106[, for a review]. We find (see Fig. 4) a strong response to all three parameters we study: concentration, shear and accretion rate, suggesting that there are additional non-negligible dependencies of b_ϕ on galaxy and halo properties besides halo mass (linear bias).

In Sec. V B, we present the first study of the parameter b_ϕ estimated from realistic galaxy samples obtained from fits to the DESI data [25]. In particular, focussing on the two most relevant tracers, LRGs and QSOs, we study the value of b_ϕ for slight variations of the main and extended (i.e., pertaining to assembly bias) HOD parameters that are still allowed by the data (see Fig. 5). We find that the uncertainty on b_ϕ is 11%, 6%, 5% for LRGs at $z = 0.5$ and $z = 0.8$ and QSOs at $z = 1.4$, respectively. We then translate these values of b_ϕ into constraints on f_{NL} from a simple Fisher analysis using linear theory. We find that the inferred

error of f_{NL}, once we take into account the uncertainty in b_ϕ, becomes 15, 8, 7 for the three tracers, respectively (see Fig. 7), which is about a factor of 1.5–2 worse compared with the case in which b_ϕ is known to a high precision. This suggests that understanding the response of galaxy formation to local-type PNG as well as the assembly bias properties of a given galaxy tracer is essential if we wish to perform unbiased, high-precision measurements of local-type PNG. For future work, we leave the exploration of the connection between observable galaxy properties and b_ϕ via hydrodynamical and N-body simulations, which would allow us to construct multitracer samples and thus greatly enhance our constraints on f_{NL}.

While we are a long way away from uncovering the subtle interplay between the physics of the primordial Universe and the observed large-scale properties of the Universe, our hope is that the AbacusPNG simulations presented in this study will aid us in this journey and provide us with some of the missing pieces needed to unravel the elusive link between fundamental cosmology and astrophysics.

The simulations used in this work are publicly available. Instructions for access and download are given at [107].

[1] N. Bartolo, E. Komatsu, S. Matarrese, and A. Riotto, Phys. Rep. **402**, 103 (2004).

[2] E. Komatsu, D. N. Spergel, and B. D. Wandelt, Astrophys. J. **634**, 14 (2005).

[3] N. Dalal, O. Doré, D. Huterer, and A. Shirokov, Phys. Rev. D **77**, 123514 (2008).

[4] P. Creminelli, G. D'Amico, M. Musso, and J. Noreña, J. Cosmol. Astropart. Phys. 11 (2011) 038.

[5] T. Tanaka and Y. Urakawa, J. Cosmol. Astropart. Phys. 05 (2011) 014.

[6] T. Baldauf, U. Seljak, and L. Senatore, J. Cosmol. Astropart. Phys. 04 (2011) 006.

[7] R. de Putter, O. Doré, and D. Green, J. Cosmol. Astropart. Phys. 10 (2015) 024.

[8] Y. Akrami *et al.* (Planck Collaboration), Astron. Astrophys. **641,** A9 (2020).

[9] N. Sailer, E. Castorina, S. Ferraro, and M. White, J. Cosmol. Astropart. Phys. 12 (2021) 049.

[10] S. Ferraro, N. Sailer, A. Slosar, and M. White, arXiv:2203.07506.

[11] A. Barreira, J. Cosmol. Astropart. Phys. 01 (2022) 033.

[12] S. Matarrese, L. Pilo, and R. Rollo, J. Cosmol. Astropart. Phys. 01 (2021) 062.

[13] V. Desjacques, U. Seljak, and I. T. Iliev, Mon. Not. R. Astron. Soc. **396,** 85 (2009).

[14] C. Wagner, L. Verde, and L. Boubekeur, J. Cosmol. Astropart. Phys. 10 (2010) 022.

[15] W. R. Coulton, F. Villaescusa-Navarro, D. Jamieson, M. Baldi, G. Jung, D. Karagiannis, M. Liguori, L. Verde, and B. D. Wandelt, Astrophys. J. **943,** 64 (2023).

[16] W. R. Coulton, F. Villaescusa-Navarro, D. Jamieson, M. Baldi, G. Jung, D. Karagiannis, M. Liguori, L. Verde, and B. D. Wandelt, Astrophys. J. **943,** 178 (2023).

[17] E. Fondi, L. Verde, F. Villaescusa-Navarro, M. Baldi, W. R. Coulton, G. Jung, D. Karagiannis, M. Liguori, A. Ravenni, and B. D. Wandelt, J. Cosmol. Astropart. Phys. 02 (2024) 048.

[18] A. G. Adame, S. Avila, V. Gonzalez-Perez, G. Yepes, M. Pellejero, C.-H. Chuang, Y. Feng, J. Garcia-Bellido, A. Knebe, and M. S. Wang, arXiv:2312.12405.

[19] A. Barreira, G. Cabass, F. Schmidt, A. Pillepich, and D. Nelson, J. Cosmol. Astropart. Phys. 12 (2020) 013.

[20] B. Hadzhiyska, S. Bose, D. Eisenstein, L. Hernquist, and D. N. Spergel, Mon. Not. R. Astron. Soc. **493,** 5506 (2020).

[21] X. Xu, I. Zehavi, and S. Contreras, Mon. Not. R. Astron. Soc. **502,** 3242 (2021).

[22] S. Yuan, L. H. Garrison, B. Hadzhiyska, S. Bose, and D. J. Eisenstein, Mon. Not. R. Astron. Soc. **510,** 3301 (2022).

[23] S. Contreras, R. E. Angulo, V. Springel, S. D. M. White, B. Hadzhiyska, L. Hernquist, R. Pakmor, R. Kannan, C. Hernández-Aguayo, M. Barrera, F. Ferlito, A. M. Delgado, S. Bose, and C. Frenk, Mon. Not. R. Astron. Soc. **524,** 2489 (2023).

[24] L. H. Garrison, D. J. Eisenstein, D. Ferrer, N. A. Maksimova, and P. A. Pinto, Mon. Not. R. Astron. Soc. **508,** 575 (2021).

[25] S. Yuan *et al.*, Mon. Not. R. Astron. Soc. **530,** 947 (2024).

[26] G. Cabass, M. M. Ivanov, O. H. E. Philcox, M. Simonović, and M. Zaldarriaga, Phys. Rev. Lett. **129,** 021301 (2022).

[27] N. A. Maksimova, L. H. Garrison, D. J. Eisenstein, B. Hadzhiyska, S. Bose, and T. P. Satterthwaite, Mon. Not. R. Astron. Soc. **508,** 4017 (2021).

[28] B. Hadzhiyska, D. Eisenstein, S. Bose, L. H. Garrison, and N. Maksimova, Mon. Not. R. Astron. Soc. **509,** 501 (2021).

[29] S. Bose, D. J. Eisenstein, B. Hadzhiyska, L. H. Garrison, and S. Yuan, Mon. Not. R. Astron. Soc. **512,** 837 (2022).

[30] A. Aghamousa *et al.* (DESI Collaboration), arXiv:1611.00036.

[31] E. Komatsu and D. N. Spergel, Phys. Rev. D **63,** 063002 (2001).

[32] L. H. Garrison, D. J. Eisenstein, D. Ferrer, M. V. Metchnik, and P. A. Pinto, Mon. Not. R. Astron. Soc. **461,** 4125 (2016).

[33] L. H. Garrison, D. J. Eisenstein, and P. A. Pinto, Mon. Not. R. Astron. Soc. **485,** 3370 (2019).

[34] L. H. Garrison, D. J. Eisenstein, and N. A. Maksimova, arXiv:2102.13140.

[35] S. Foreman, W. Coulton, F. Villaescusa-Navarro, and A. Barreira, Mon. Not. R. Astron. Soc. **498,** 2887 (2020).

[36] X. Chen, Adv. Astron. **2010,** 638979 (2010).

[37] A. Achúcarro *et al.*, arXiv:2203.08128.

[38] J. Maldacena, J. High Energy Phys. 05 (2003) 013.

[39] P. Creminelli and M. Zaldarriaga, J. Cosmol. Astropart. Phys. 10 (2004) 006.

[40] D. H. Lyth and D. Wands, Phys. Lett. B **524,** 5 (2002).

[41] G. Dvali, A. Gruzinov, and M. Zaldarriaga, Phys. Rev. D **69,** 023505 (2004).

[42] J.-L. Lehners, Adv. Astron. **2010,** 903907 (2010).

[43] P. McDonald, Phys. Rev. D **78,** 123519 (2008).

[44] V. Assassi, D. Baumann, E. Pajer, Y. Welling, and D. v. d. Woude, J. Cosmol. Astropart. Phys. 11 (2015) 024.

[45] V. Assassi, D. Baumann, and F. Schmidt, J. Cosmol. Astropart. Phys. 12 (2015) 043.

[46] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, J. Cosmol. Astropart. Phys. 08 (2008) 031.

[47] S. Matarrese and L. Verde, Astrophys. J. Lett. **677,** L77 (2008).

[48] N. Afshordi and A. J. Tolley, Phys. Rev. D **78,** 123507 (2008).

[49] P. Valageas, Astron. Astrophys. **514,** A46 (2010).

[50] T. Matsubara, Phys. Rev. D **86,** 063518 (2012).

[51] S. Ferraro, K. M. Smith, D. Green, and D. Baumann, Mon. Not. R. Astron. Soc. **435,** 934 (2013).

[52] R. Scoccimarro, L. Hui, M. Manera, and K. C. Chan, Phys. Rev. D **85,** 083002 (2012).

[53] M. Biagetti, T. Lazeyras, T. Baldauf, V. Desjacques, and F. Schmidt, Mon. Not. R. Astron. Soc. **468,** 3277 (2017).

[54] M. Grossi, L. Verde, C. Carbone, K. Dolag, E. Branchini, F. Iannuzzi, S. Matarrese, and L. Moscardini, Mon. Not. R. Astron. Soc. **398,** 321 (2009).

[55] A. Pillepich, C. Porciani, and O. Hahn, Mon. Not. R. Astron. Soc. **402,** 191 (2010).

[56] B. A. Reid, L. Verde, K. Dolag, S. Matarrese, and L. Moscardini, J. Cosmol. Astropart. Phys. 07 (2010) 013.

[57] N. Hamaus, U. Seljak, and V. Desjacques, Phys. Rev. D **84,** 083509 (2011).

[58] C. Wagner and L. Verde, J. Cosmol. Astropart. Phys. 03 (2012) 002.

[59] T. Baldauf, U. Seljak, L. Senatore, and M. Zaldarriaga, J. Cosmol. Astropart. Phys. 09 (2016) 007.

[60] J.-Q. Xia, C. Baccigalupi, S. Matarrese, L. Verde, and M. Viel, J. Cosmol. Astropart. Phys. 08 (2011) 033.

[61] A. J. Ross *et al.*, Mon. Not. R. Astron. Soc. **428,** 1116 (2013).

[62] T. Giannantonio, A. J. Ross, W. J. Percival, R. Crittenden, D. Bacher, M. Kilbinger, R. Nichol, and J. Weller, Phys. Rev. D **89,** 023511 (2014).

[63] B. Leistedt, H. V. Peiris, and N. Roth, Phys. Rev. Lett. **113,** 221301 (2014).

[64] T. Giannantonio and W. J. Percival, Mon. Not. R. Astron. Soc. **441,** L16 (2014).

[65] S. Ho *et al.*, J. Cosmol. Astropart. Phys. 05 (2015) 040.

[66] E. Castorina, N. Hand, U. Seljak, F. Beutler, C.-H. Chuang, C. Zhao, H. Gil-Marín, W. J. Percival, A. J. Ross, P. D. Choi, K. Dawson, A. de la Macorra, G. Rossi, R. Ruggeri, D. Schneider, and G.-B. Zhao, J. Cosmol. Astropart. Phys. 09 (2019) 010.

[67] C. Carbone, L. Verde, and S. Matarrese, Astrophys. J. Lett. **684,** L1 (2008).

[68] T. Giannantonio, C. Porciani, J. Carron, A. Amara, and A. Pillepich, Mon. Not. R. Astron. Soc. **422,** 2854 (2012).

[69] O. Doré *et al.*, arXiv:1412.4872.

[70] M. Alvarez *et al.*, arXiv:1412.4671.

[71] A. Raccanelli, O. Doré, D. J. Bacon, R. Maartens, M. G. Santos, S. Camera, T. M. Davis, M. J. Drinkwater, M. Jarvis, R. Norris, and D. Parkinson, J. Cosmol. Astropart. Phys. 01 (2015) 042.

[72] D. Alonso and P. G. Ferreira, Phys. Rev. D **92,** 063525 (2015).

[73] S. Camera, M. G. Santos, and R. Maartens, Mon. Not. R. Astron. Soc. **448,** 1035 (2015).

[74] R. de Putter and O. Doré, Phys. Rev. D **95,** 123513 (2017).

[75] A. Raccanelli, M. Shiraishi, N. Bartolo, D. Bertacca, M. Liguori, S. Matarrese, R. P. Norris, and D. Parkinson, Phys. Dark Universe **15,** 35 (2017).

[76] N. Kaiser, Astrophys. J. Lett. **284,** L9 (1984).

[77] J. M. Bardeen, J. R. Bond, N. Kaiser, and A. S. Szalay, Astrophys. J. **304,** 15 (1986).

[78] U. Giri, M. Münchmeyer, and K. M. Smith, Phys. Rev. D **107,** L061301 (2023).

[79] J. L. Tinker, B. E. Robertson, A. V. Kravtsov, A. Klypin, M. S. Warren, G. Yepes, and S. Gottlöber, Astrophys. J. **724,** 878 (2010).

[80] A. Barreira, J. Cosmol. Astropart. Phys. 12 (2020) 031.

[81] Y.-Y. Mao, A. R. Zentner, and R. H. Wechsler, Mon. Not. R. Astron. Soc. **474,** 5143 (2018).

[82] J. F. Navarro, C. S. Frenk, and S. D. M. White, Astrophys. J. **490,** 493 (1997).

[83] R. H. Wechsler, J. S. Bullock, J. R. Primack, A. V. Kravtsov, and A. Dekel, Astrophys. J. **568,** 52 (2002).

[84] A. D. Ludlow, J. F. Navarro, R. E. Angulo, M. Boylan-Kolchin, V. Springel, C. Frenk, and S. D. M. White, Mon. Not. R. Astron. Soc. **441,** 378 (2014).

[85] A. D. Ludlow, S. Bose, R. E. Angulo, L. Wang, W. A. Hellwing, J. F. Navarro, S. Cole, and C. S. Frenk, Mon. Not. R. Astron. Soc. **460,** 1214 (2016).

[86] K. Wang, Y.-Y. Mao, A. R. Zentner, J. U. Lange, F. C. van den Bosch, and R. H. Wechsler, Mon. Not. R. Astron. Soc. **498,** 4450 (2020).

[87] J. S. Bullock, T. S. Kolatt, Y. Sigad, R. S. Somerville, A. V. Kravtsov, A. A. Klypin, J. R. Primack, and A. Dekel, Mon. Not. R. Astron. Soc. **321,** 559 (2001).

[88] B. Diemer and A. V. Kravtsov, Astrophys. J. **799,** 108 (2015).

[89] A. A. Dutton and A. V. Macciò, Mon. Not. R. Astron. Soc. **441,** 3359 (2014).

[90] B. Hadzhiyska, D. Eisenstein, S. Bose, L. H. Garrison, and N. Maksimova, Mon. Not. R. Astron. Soc. **509,** 501 (2022).

[91] G. L. Bryan and M. L. Norman, Astrophys. J. **495,** 80 (1998).

[92] B. Hadzhiyska, D. Eisenstein, L. Hernquist, R. Pakmor, S. Bose, A. M. Delgado, S. Contreras, R. Kannan, S. D. M. White, V. Springel, C. Frenk, C. Hernández-Aguayo, F. Ferlito, and Monica Barrera, Mon. Not. R. Astron. Soc. **524,** 2507 (2023).

[93] S. Yuan, B. Hadzhiyska, S. Bose, D. J. Eisenstein, and H. Guo, Mon. Not. R. Astron. Soc. **502,** 3582 (2021).

[94] B. Hadzhiyska, S. Bose, D. Eisenstein, and L. Hernquist, Mon. Not. R. Astron. Soc. **501,** 1603 (2021).

[95] J. Chaves-Montero, R. E. Angulo, and S. Contreras, Mon. Not. R. Astron. Soc. **521,** 937 (2023).

[96] Z. Zheng, A. A. Berlind, D. H. Weinberg, A. J. Benson, C. M. Baugh, S. Cole, R. Davé, C. S. Frenk, N. Katz, and C. G. Lacey, Astrophys. J. **633,** 791 (2005).

[97] H. Guo, Z. Zheng, I. Zehavi, K. Dawson, R. A. Skibba, J. L. Tinker, D. H. Weinberg, M. White, and D. P. Schneider, Mon. Not. R. Astron. Soc. **446,** 578 (2015).

[98] S. D. Landy and A. S. Szalay, Astrophys. J. **412,** 64 (1993).

[99] M. Tegmark, A. N. Taylor, and A. F. Heavens, Astrophys. J. **480,** 22 (1997).

[100] A. C. Aitken and H. Silverstone, Proc. R. Soc. Edinburgh, Sect. A **61,** 186 (1942).

[101] B. Hadzhiyska *et al.*, Open J. Astrophys. **6,** 38 (2023).

[102] J. Carron, Astron. Astrophys. **551,** A88 (2013).

[103] O. Alves (The DESI Collaboration) (to be published).

[104] J. Martin, Astrophys. Space Sci. Proc. **45,** 41 (2016).

[105] P. D. Meerburg, D. Green, R. Flauger, B. Wallisch, M. C. D. Marsh, E. Pajer, G. Goon, C. Dvorkin, A. M. Dizgah, D. Baumann, G. L. Pimentel, S. Foreman, E. Silverstein, E. Chisari, B. Wandelt, M. Loverde, and A. Slosar, Bull. Am. Astron. Soc. **51,** 107 (2019).

[106] R. H. Wechsler and J. L. Tinker, Annu. Rev. Astron. Astrophys. **56,** 435 (2018).

[107] https://abacussummit.readthedocs.io/en/latest/data-access.html.